# Exhibit 2

UNITED STATES PATENT AND TRADEMARK OFFICE

---

BEFORE THE PATENT TRIAL AND APPEAL BOARD

---

MEDIATEK, INC. and
MEDIATEK USA, INC.,

Petitioners,

v.

REDSTONE LOGICS LLC,

Patent Owners.

---

IPR2025-00085
U.S. Patent No. 8,549,339

---

**DECLARATION OF R. JACOB BAKER, PH.D., P.E.
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW
OF U.S. PATENT NO. 8,549,339**

# I. INTRODUCTION

1. My name is R. Jacob Baker Ph.D., P.E., and I am an Emeritus Professor of Electrical and Computer Engineering at the University of Nevada, Las Vegas ("UNLV"). I have prepared this declaration as an expert witness on behalf of Petitioners MediaTek, Inc., and MediaTek USA, Inc. In this declaration, I provide my opinions regarding the validity of claims 1-6, 8-11, 14 and 21 (the "Challenged Claims") of U.S. Patent No. 8,549,339 ("the '339 Patent") (Ex[1001]). I also provide herein the technical bases for these opinions, as appropriate.

2. This declaration contains statements of my opinions formed to date, and the bases and rationale for these opinions. I may offer additional opinions based on further review of materials in this case, including opinions and/or testimony of other expert witnesses.

3. For my efforts in connection with the preparation of this declaration, I have been compensated at my usual and customary rate for this type of consulting activity. My compensation is in no way contingent on the results of these or any other proceedings related to the '339 Patent.

## A. Qualifications / Professional Background

4. My qualifications generally are set forth in my Curriculum Vitae, which is included as Ex[1004]. Ex[1004] also includes a list of the publications I have authored and a list of the other cases in which I have testified during the last four

years.

5. I have been working as an engineer since 1985, and I have been teaching Electrical and Computer Engineering courses since 1991.

6. I received B.S. and M.S. degrees in Electrical Engineering from UNLV in 1986 and 1988, respectively. I received my Ph.D. in Electrical Engineering from the University of Nevada, Reno ("UNR"), in 1993.

7. My doctoral research, culminating in the award of a Ph.D., investigated the use of power MOSFETs in the design of very high peak power, and high-speed, instrumentation. I developed techniques to reliably stack power MOSFETs to switch higher voltages, that is, greater than 1,000 V and 100 Amps of current with nanosecond switching times. This work was reported in the paper entitled "Transformerless Capacitive Coupling of Gate Signals for Series Operation of Power MOSFET Devices," published in the IEEE Transactions on Power Electronics. The paper received the Best Paper Award in 2000.

**B. Industry Experience**

8. From 1985 to 1993, I worked for EG&G Energy Measurements and the Lawrence Livermore National Laboratory designing nuclear diagnostic instrumentation for underground nuclear weapon tests at the Nevada test site. During this time, I designed, and oversaw the fabrication of, over 30 electronic and electro-optic instruments, including high-speed cable and fiber-optic receiver/transmitters,

PLLs, frame and bit-syncs, data converters, streak-camera sweep circuits, Pockel's cell drivers, micro-channel plate gating circuits, charging circuits for battery backup of equipment for recording test data, and analog oscilloscope electronics.

9. My work during this time, as one example, had a direct impact on my doctoral research work using power MOSFETs, subsequent publishing efforts, and industry designs. In addition to the 2000 Best Paper Award from the IEEE Power Electronics Society, I published several other papers in related areas while working in industry. I hold a patent, U.S. Patent No. 5,874,830, in the area of power supply design, titled, "Adaptively Biased Voltage Regulator and Operating Method," which was issued on February 23, 1999. I have designed dozens of linear and switching power supplies for commercial products and scientific instrumentation.

10. I am a licensed Professional Engineer and have extensive industry experience in circuit design, fabrication, and manufacture of Dynamic Random Access Memory (DRAM) semiconductor integrated circuit chips, Phase-Change Random Access Memory (PCRAM) chips, and CMOS Image Sensors (CISs) at Micron Technology, Inc. ("MTI") in Boise, Idaho. I spent considerable time working on the development of flash memory chips while at MTI. My efforts resulted in more than a dozen patents relating to flash memory. One of my projects at MTI included the development, design, and testing of circuit design techniques for a multi-level cell (MLC) flash memory using signal processing. This effort resulted in higher-

density memories for use in solid-state drives and flash memory cards having an ATA interface that are ubiquitous in consumer electronics, including cameras and data storage systems. Further, the use of higher-density memory can result in fewer changes in the flash translation layer for logical-to-physical addressing, less need for garbage collection, and larger data segments that can improve a computing system's performance. Another project I worked on at MTI focused on the design of buffers for high-speed double-data rate DRAM, which resulted in around 10 U.S. patents in buffer design. Among many other experiences, I led the development of the delay locked loop (DLL) in the late 1990s so that MTI DRAM products could transition to the DDR memory protocol, used in mobile and non-mobile (server, desktop, cell phones, tablets, etc.) computing systems as main computer memory, for addressing and controlling accesses to memory via interprocess communications (IPC) with the memory controller (MC). I provided technical assistance with MTI's acquisition of Photobit during 2001 and 2002, including transitioning the manufacture of CIS products into MTI's process technology. Further, I did consulting work at Sun Microsystems and then Oracle on the design of memory modules during 2009 and 2010. This work entailed the design of low-power, high-speed, and wide interconnection methods with the goal of transmitting data to/from the memory module and the MC at higher speeds.

11. I have extensive experience in the development of instrumentation and

commercial products in a multitude of areas including: integrated electrical/biological circuits and systems, array (memory, imagers, and displays) circuit design, CMOS analog and digital circuit design, diagnostic electrical and electro-optic instrumentation for scientific research, CAD tool development and online tutorials, low-power interconnect and packaging techniques, design of communication/interface circuits (to meet commercial standards such as USB, firewire, DDR, PCIe, SPI, etc.), circuit design for the use and storage of renewable energy, and power electronics. For example, a part of my research at Boise State, for many years, focused on the use of Thru-Silicon-Vias (TSVs), aka Thru-Wafer Vias (TWVs), for high-density 3D packaging. These packaging techniques were utilized in the memory module development work I did with Sun Microsystems and Oracle. As another example, I designed circuitry for use in implementing Universal Serial Bus (USB) interface circuits while I did consulting at Tower Semiconductor. I designed PCI communication circuits for IPC between a Graphics Processor Unit (GPU) and memory while consulting for Rendition, Inc. From 1994 to 1996, I worked on the design of displays at MTI. This work was at a time when cathode ray tubes (CRTs) were still the dominant type of display. Flat panel displays were being developed with the hope of replacing CRTs in the consumer market. I worked on flat panel displays which resulted in five patents: U.S. Patent Nos. 5,598,156, 5,638,085, 5,818,365, 5,894,293, and 5,909,201. I worked on the design of the

pixels, both active and passive, as well as the supporting electronics for processing video signals. I was involved with the evaluation of display technologies including liquid crystal displays (LCDs), light-emitting diodes, plasma displays, and organic light emitting diode (OLED) displays; the display technologies that were looking to displace CRTs in the consumer market. I was also involved with the packaging of the displays including the vacuum sealing and deposition of the phosphors for light wavelength conversion. I also taught display design as a topic in my courses and did display design consulting again in industry for Cirque in 2013.

12. My current research work is focused in part on the design of integrated circuits for wireless sensing using LIDAR (LIght Detection And Ranging). I have worked with several companies in the development of these circuits and systems including Freedom Photonics, Aerius Photonics, and FLIR. In the early 1990s, I worked on wireless systems for wideband impulse radar while at Lawrence Livermore Laboratory. Further, part of my research for several years focused on the digitization of IQ channels using delta-sigma modulation. The knowledge and experience gained from this effort are reflected in my textbook CMOS Mixed-Signal Circuit Design and a presentation, which I have presented at several universities and companies, http://cmosedu.com/jbaker/papers/talks/BP_DSM_talk.pdf.

### C. Academic Experience

13. I was an adjunct faculty member in the Electrical Engineering

departments of UNLV and UNR from 1991-1993. From 1993 to 2000, I served on the faculty at the University of Idaho as an Assistant Professor and then as a tenured Associate Professor of Electrical Engineering. In 2000, I joined a new Electrical and Computer Engineering program at Boise State University ("BSU"), where I served as department chair from 2004 to 2007. At BSU, I helped establish graduate programs in Electrical and Computer Engineering including, in 2006, the university's second Ph.D. degree. In 2012, I re-joined the faculty at UNLV. Over the course of my career as a professor, I have advised more than 100 masters and doctoral students.

14. I have been recognized for my contributions as an educator in the field. While at BSU, I received the President's Research and Scholarship Award (2005), Honored Faculty Member recognition (2003), and Outstanding Department of Electrical Engineering Faculty recognition (2001). In 2007, I received the Frederick Emmons Terman Award (the "Father of Silicon Valley"). The Terman Award is bestowed annually upon an outstanding young electrical/computer engineering educator in recognition of the educator's contributions to the profession. In 2011, I received the IEEE Circuits and Systems Education Award. I received the Tau Beta Pi Outstanding Electrical and Computer Engineering Professor Award every year it was awarded while I have been back at UNLV.

15. I have authored several books and papers in the electrical and computer

engineering area. My published books include CMOS Circuit Design, Layout, and Simulation (Baker, R.J., Wiley-IEEE, ISBN: 9781119481515 (4th ed., 2019)) and CMOS Mixed-Signal Circuit Design (Baker, R.J., Wiley-IEEE, ISBN: 9780470290262 (2nd ed., 2009) and ISBN: 9780471227540 (1st ed., 2002)). I co-authored DRAM Circuit Design: Fundamental and High-Speed Topics (Keeth, B., Baker, R.J., Johnson, B., and Lin, F., Wiley-IEEE, ISBN: 9780470184752 (2008)), DRAM Circuit Design: A Tutorial (Keeth, B. and Baker, R.J., Wiley-IEEE, ISBN: 0-7803-6014-1 (2001)), and CMOS Circuit Design, Layout and Simulation (Baker, R.J., Li, H.W., and Boyce, D.E., Wiley - IEEE, ISBN: 9780780334168 (1998)). I contributed as an editor and co-author on several other electrical and computer engineering books.

**D. Other Relevant Experience**

16. I have performed technical and expert witness consulting for more than 200 companies and their subsidiaries and given more than 50 invited talks at conferences, companies, and universities. Further, I am the author or co-author of more than 100 papers and presentations in the areas of electrical and computer engineering design, fabrication, and packaging.

17. I currently serve, or have served, as a volunteer on the IEEE Press Editorial Board (1999-2004); as editor for the Wiley-IEEE Press Book Series on Microelectronic Systems (2010-2018); as the Technical Program Chair of the 2015

IEEE 58th International Midwest Symposium on Circuits and Systems (MWSCAS 2015); on the IEEE Solid-State Circuits Society (SSCS) Administrative Committee (2011-2016); as a Distinguished Lecturer for the SSCS (2012-2015); the Technology Editor (2012-2014) and Editor-in-Chief (2015-2020) for IEEE Solid-State Circuits Magazine; IEEE Kirchhoff Award Committee (2020-2023); and advisor for the student branch of the IEEE at UNLV (2013-present). These meetings, groups, and publications are intended to allow researchers to share and coordinate research. My active participation in these meetings, groups, and publications allowed me to see what other researchers in the field have been doing.

18. In addition to the above, I am an IEEE Fellow for contributions to semiconductor memory design and a member of the honor societies Eta Kappa Nu and Tau Beta Pi.

## II. SUMMARY AND OPINIONS OF THE '339 PATENT

### A. Overview of the '339 Patent

19. The '339 Patent purports to teach techniques for handling communication between processor cores of a multi-core processor. Ex[1001] at Abstract. The objective of the '339 Patent is provided below:

> ***Embodiments of the disclosure generally set forth techniques for handling communication between processor cores.*** Some example multi-core processors include a first set of processor cores in a first region of the multi-core processor configured to dynamically receive a first supply voltage and a first clock signal, a second set of processor cores in a second region of the multi-core processor configured to

> dynamically receive a second supply voltage and a second clock signal, and an interface block coupled to the first set of processor cores and the second set of processor cores, wherein the interface block is configured to facilitate communications between the first set of processor cores and the second set of processor cores.

*Id.* at Abstract.

20. Figure 1 shows four sets of processor cores, with the first set (highlighted in green) and the second set (highlighted in red) of processor cores highlighted.[1]

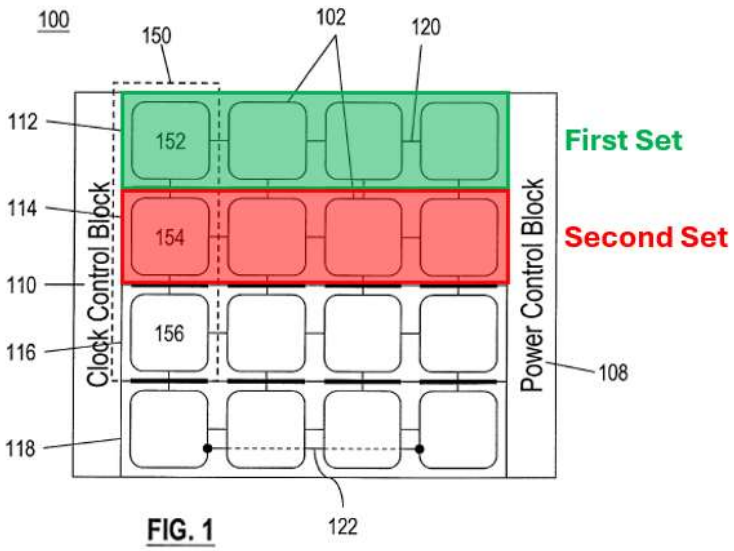


FIG. 1

21. Figures 2 and 3 show an interface block (highlighted in blue) coupled to the first set of processor cores (highlighted in green) and also coupled to the second set of processor cores (highlighted in red):

[1] Throughout, I have added emphasis and annotations unless otherwise specified.

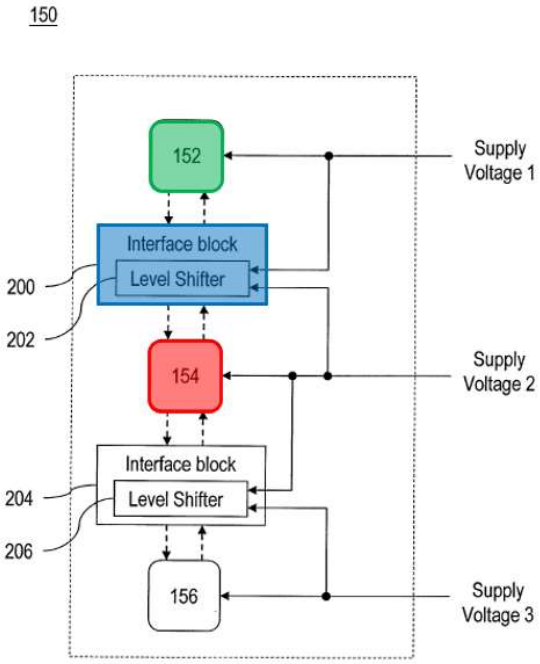


FIG. 2

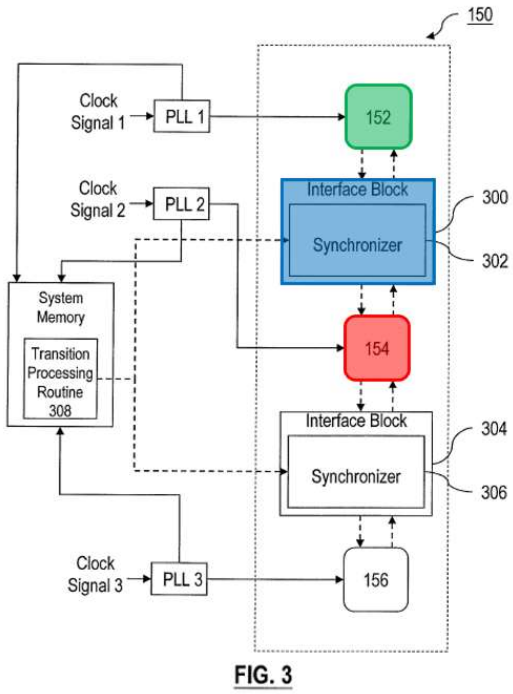


FIG. 3

22. The '339 Patent discloses that for power consumption management, dynamic voltage supply and clock speed control is utilized for the first set and the second set of processor cores, such that the multi-core processor may operate at high power and high clock frequency when needed and at low power when the computing requirements are reduced. *Id.* at 1:10-14. In this regard, the object of the '339 Patent is as follows:

> A multi-core processor includes two or more independent processor cores arranged in an array. Each processor core in a conventional multi-core processor generally shares the same supply voltage and clock signal to simplify the interfaces between the processor cores. ***For power consumption management, dynamic supply voltage and clock speed control may be utilized, so that a multi-core processor may operate at high power and high clock frequency when needed and at low power when the computing requirements are reduced.***

*Id.* at 1:6-14.

23. The power profiles for the first set and the second set of processor cores are independent of each other. Similarly, the clock signals of the first set and the second set of processor cores are independent of each other. *Id.* at 2:25-31. Further elaboration is provided below:

> For example, the multi-core processor 100 may be divided into stripes 112, 114, 116, and 118. ***Each stripe may be associated with an independent power profile. For example, the stripe 112 may be powered by a supply voltage received from a power control block 108 and/or may be associated with an independent clock domain defined by a clock signal received from a clock control block 110.***

*Id.* at 2:25-31.

> ***The power profile associated with a stripe may be determined based on the computational requirements of the tasks assigned to the processor cores in the stripe***.

*Id.* at 2:41-43.

24. As such, facilitating communication between two processor cores is necessary and is performed by interface blocks. *Id.* at 3:21-23. Further disclosures of interface blocks, such as having a level shifter or a synchronizer, are provided below:

> ***The processor core 152 may be powered by a supply voltage 1 and coupled to an interface block 200 having a level shifter 202***; the processor core 154 may be powered by a supply voltage 2 and coupled to the same interface block 200; and ***the processor core 156 may be powered by a supply voltage 3 and coupled to an interface block 204 having a level shifter 206.***

*Id.* at 3:30-36.

> ***The processor core 152 may be driven by a clock signal 1 and coupled to an interface block 300 having a synchronizer 302***; the processor core 154 may be driven by a clock signal 2 and coupled to the same interface block 300; and ***the processor core 156 may be driven by a clock signal 3 and coupled to an interface block 304 having a synchronizer 306***.

*Id.* at 4:4-10.

### B. Priority Date for the '339 Patent

25. I have reviewed the prosecution history of the '339 Patent. The Challenged Claims issued from U.S. Patent Appl. No. 12/713,220 (the "'220 Application") and do not purport to claim domestic priority or foreign priority to any other application. The earliest possible priority for the '339 Patent is therefore February 26, 2010. I have been asked to assume this earliest date as the priority date. All references asserted herein are prior art.

## III. LEGAL BASES OF OPINIONS

### A. Level of Ordinary Skill in the Art

26. I have been informed that both the '339 Patent and the prior art are to be understood from the perspective of a hypothetical person having ordinary skill in the art at the time of the purported inventions claimed in the '339 Patent (a "PHOSITA"). I have been asked to treat the earliest possible priority date of February 26, 2010, as the time of the purported inventions claimed in the '339 Patent for purposes of this proceeding.

27. I understand that the factors that may be considered to determine the

level of ordinary skill in the art include: (a) the type of problems encountered in the art; (b) prior art solutions to those problems; (c) the complexity or sophistication of the technology in the art; (d) the education level and experience of active workers in the field of art; and (e) the pace of change, development, and innovation in the field of art. I have been informed that not every factor would be present in every case, and that a review of the prior art would shed light on factors (a), (b), (c), and (e). I have also been informed that a PHOSITA is further presumed to have been aware of all relevant prior art.

28. As discussed in further detail below, the ’339 Patent is directed to communication between processor cores of a multi-core processor. The claims incorporate common elements including a multi-core processor, supply voltages, clock signals, and interface blocks. The relevant prior art relates to sets of processor cores and the communication between them. I am familiar with the state of the art in this field at the time of the purported inventions claimed in the ’339 Patent.

29. In my opinion, a PHOSITA would have had a minimum of a bachelor’s degree in electrical engineering, computer engineering, computer science, or a similar field, and at least two years of industry or academic experience designing or analyzing electronic circuits, semiconductors, processors, or power management, and related firmware and software, or the equivalent. The more education one has (e.g., additional studies, coursework, or degrees), the less experience needed to attain

an ordinary level of skill. Likewise, more extensive experience might substitute for some educational requirements.

30. In view of my educational background and my decades of academic and professional experience (as discussed above), I was a person of more than the ordinary level of skill in the art in the early 2010 timeframe. My opinions expressed in this declaration were formed, however, from the perspective of a person of ordinary skill in the art as of February 26, 2010. I was, and am, well acquainted with the level of knowledge of a PHOSITA under this definition. As noted above, I have worked with many individuals during my time in the industry and taught many who would have met the definition of a PHOSITA. Thus, I am comfortable applying this viewpoint in assessing the '339 Patent and the state of the art in the relevant timeframe, including the prior art references I discuss below.

31. I have also considered whether my opinions expressed in this declaration would change if the level of ordinary skill in the art was varied by some minor amount of time or varied somewhat with respect to subject matter. My opinions would not change in light of such minor variations.

**B. Claim Interpretation**

32. I have been informed that patent claims are to be interpreted from the perspective of PHOSITA. I have also been informed that the words of a patent claim are generally given their ordinary and customary meaning to a PHOSITA in view of

the context of the claims and the intrinsic evidence, *i.e.*, the patent specification and the prosecution history of the '339 Patent. I have read and analyzed the claims of the '339 Patent from this perspective.

33. I have been further informed that extrinsic evidence, *e.g.*, prior art publications and expert testimony, may be considered to understand a patent claim, but cannot be used to contradict an interpretation thereof which is unambiguous in view of the intrinsic evidence.

### C. Obviousness

34. As discussed above, the '339 Patent has an effective filing date of February 26, 2010, and I have been asked to treat this date as the time of the purported inventions claimed in the '339 Patent.

35. I have been informed that a claimed invention is obvious — and thus unpatentable — if the differences between the claimed invention and the prior art are such that the claimed invention as a whole would have been obvious to a PHOSITA at the time of the claimed invention.

36. I have also been informed that the determination of whether a claimed invention is obvious in view of the prior art cannot rely on hindsight, *i.e.*, using the patent claim as a roadmap or template to select components from the prior art to reconstruct the claimed invention.

37. I have been informed that a PHOSITA is presumed to have been aware

of all relevant prior art. I have also been informed that a PHOSITA is someone of ordinary creativity in the field who is presumed to have understood and considered design needs and/or market forces which would have motivated or prompted them to seek variations of a product having utility in their field.

38. I have been informed that the determination of whether a claimed invention was obvious in view of the prior art is a legal question that involves an objective analysis of these factual factors:

- The scope and content of the prior art;
- The differences between the prior art and the claimed invention;
- The level of ordinary skill in the pertinent art; and
- Any secondary consideration(s) having a nexus to the claimed invention, for example: [i] its commercial success (or the lack thereof), [ii] whether the claimed invention satisfied a long felt but unmet need, [iii] unexpected results achieved by the claimed invention, or [iv] skepticism or failure by others.

39. I have also been informed that the above factors define an expansive and flexible framework to determine whether a claimed invention would have been obvious in view of the prior art. For example:

40. When a prior art reference was available in one field, design incentives and other market forces could prompt variations of the reference by a PHOSITA,

either in the same field or a different field. If the PHOSITA could implement a predictable variation that meets the elements of a patent claim, then the claimed invention was obvious.

41. Likewise, if a technique had been used to improve one device in the prior art, and a PHOSITA would recognize that the technique could also be used to improve similar devices in the same way, using that technique was obvious unless its actual application was beyond the capability of the PHOSITA.

42. The combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results. If an invention is no more than the predictable use of prior art elements according to their established functions, the combination was obvious unless its implementation was beyond the capability of the PHOSITA.

43. To determine whether there was an apparent reason to combine elements known in the prior art as claimed, it is often necessary to examine the interrelated teachings of multiple prior art references, the effects of demands known to the field or present in the market, and the background knowledge of a PHOSITA. This determination is not a rigid test. Rather, it must consider the inferences and creative steps that a PHOSITA would employ, because the PHOSITA is not an automaton.

44. Common sense teaches that familiar items may have obvious uses

beyond their primary purposes, and in many cases a PHOSITA would have been able to fit the teachings of multiple patents together like pieces of a puzzle.

45. A motivation to combine the teachings of multiple prior art references may be suggested in one or more of the references, or found in the nature of the problem to be solved, in the trends or direction of the field and/or market, or come from the common sense and ordinary skill of the PHOSITA.

46. When there was a design need or market pressure to solve a problem and there were a finite number of identified, predictable solutions, a PHOSITA would have good reason to pursue the known options within his or her technical grasp. If this would have led to reasonably anticipated success, it was likely the product not of innovation but of ordinary skill and common sense, in which case the fact that a combination of known elements was obvious to try would show that the claimed invention was obvious.

47. Certain aspects of the purported inventions claimed in the ’366 Patent, for example, fans and fasteners, pertain to the mechanical arts. I have been informed that the mechanical arts have long been recognized as a predictable art, one where the combination of elements known in the prior art typically yields predictable results.

**D. Burden of Proof**

48. I have been informed that a petitioner must demonstrate unpatentability

by a preponderance of the evidence. I understand preponderance means more likely than not in this context.

## IV. CHALLENGED CLAIMS OF THE '339 PATENT

49. The following claims appear in the '339 Patent:

| | |
|---|---|
| 1[pre] | A multi-core processor, comprising: |
| 1[a1] | a first set of processor cores of the multi-core processor, |
| 1[a2] | wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage and a first output clock signal of a first phase lock loop (PLL) having a first clock signal as input; |
| 1[b1] | a second set of processor cores of the multi-core processor, |
| 1[b2] | wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage and a second output clock signal of a second PLL having a second clock signal as input, |
| 1[b3] | wherein the first supply voltage is independent from the second supply voltage, and |
| 1[b4] | the first clock signal is independent from the second clock signal; and |
| 1[c1] | an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores, |
| 1[c2] | wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores. |
| 2[pre] | The multi-core processor of claim 1, |
| 2[a] | wherein the interface block further comprises a first level shifter that is referenced to the second supply voltage and adapted to translate first logic levels associated with the first set of processor cores to second logic levels associated with the second set of processor cores for a first signal traveling from the first set of processor cores to the second set of processor cores. |
| 3[pre] | The multi-core processor of claim 1, |

| | |
|---|---|
| 3[a] | wherein the interface block further comprises a second level shifter that is referenced to the first supply voltage and adapted to translate second logic levels associated with the second set of processor cores to first logic levels associated with the first set of processor cores for a second signal traveling from the second set of processor cores to the first set of processor cores. |
| 4[pre] | The multi-core processor of claim 1, |
| 4[a] | wherein the interface block further comprises a synchronizer configured to synchronize the first clock signal and the second clock signal for communication between one or more processor cores of the first set of processor cores and one or more processor cores of the second set of processor cores. |
| 5[pre] | The multi-core processor of claim 1, |
| 5[a] | wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor. |
| 6[pre] | The multi-core processor of claim 5, |
| 6[a1] | wherein the first set of processor cores is adjacent to the second set of processor cores, and |
| 6[a2] | the one or more control blocks are configured to select the first supply voltage and the second supply voltage to maintain a differential relationship between the first supply voltage and the second supply voltage. |
| 8[pre] | The multi-core processor of claim 1, |
| 8[a] | wherein the first set of processor cores are located in a first region of the multi-core processor, and the second set of processor cores are located in a second region of the multi-core processor. |
| 9[pre] | The multi-core processor of claim 8, |
| 9[a] | wherein the first region and the second region are overlapping regions of the multi-core processor. |
| 10[pre] | The multi-core processor of claim 8, |
| 10[a] | wherein the first region and the second region are non-overlapping regions of the multi-core processor. |
| 11[pre] | The multi-core processor of claim 8, |
| 11[a] | wherein the first region corresponds to a first row of the multi-core processor, and wherein the second region corresponds to a second row of the multi-core processor. |
| 14[pre] | The multi-core processor of claim 1, |

| | |
|---|---|
| 14[a] | wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a common region that is substantially central to the first set of processor cores and the second set of processor cores. |
| 21[pre] | A multi-core processor, comprising: |
| 21[a1] | a first set of processor cores of the multi-core processor, |
| 21[a2] | wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage from a power control block and a first output clock signal from a first phase lock loop (PLL) having a first clock signal as input in a clock control block; |
| 21[b1] | a second set of processor cores of the multi-core processor, |
| 21[b2] | wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage from the power control block and a second output clock signal from a second PLL having a second clock signal as input in the clock control block, |
| 21[b3] | wherein the first supply voltage is independent from the second supply voltage, and |
| 21[b4] | the first clock signal is independent from the second clock signal; and |
| 21[c1] | an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores, |
| 21[c2] | wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores. |

## V. OVERVIEW OF PRIOR ART REFERENCES

### A. U.S. Patent Application Publication No. 2009/0158078 A1 to Knoth ("Knoth") (Ex[1005])

50. U.S. Patent Application. No. 2009/0158078 ("Knoth") is a published U.S. patent application filed on December 12, 2007, and published on June 18, 2009. Ex[1005] at 1. Its title is "Clock Ratio Controller for Dynamic Voltage and

Frequency Scaled Digital Systems, and Applications Thereof." Knoth is therefore prior art to the '339 Patent under at least 35 U.S.C. §§ 102(a) and 102(e)(1).

51. Knoth provides a clock ratio controller for dynamic voltage and frequency scaling without a loss of synchronization between system elements. *Id.* at [0003]. *See* below *id.* at FIGS. 1A and 2A.

> ***The present invention provides a clock ratio controller for dynamic voltage and frequency scaled digital systems, and applications thereof.*** In an embodiment, ***the clock ratio controller is used, for example, to adjust the frequency of clock signals of a digital system without a loss of synchronization between system elements.***

*Id.* at [0003].

> In an embodiment, power management unit 102 provides individual frequency management (FM) signals 111a-n to clock ratio controllers 106a-n and individual voltage management (VM) signals 113a-n to voltage controllers 107a-n. These signals are also referred to herein as power management signals. A FM signal 111o is also provided to prescaler 104. ***These FM signals are used, for example, to initiate frequency adjustments. The VM signals are used, for example, to initiate voltage adjustments. FM signals 111a-n can be used to individually control/scale the frequency of each processor core 110a-n. VM signals 113a-n can be used to individually control/scale the voltage of each processor core 110a-n.*** FM signal 111o can be used to control/scale the frequency of all of the components of digital system 100.

*Id.* at [0025].

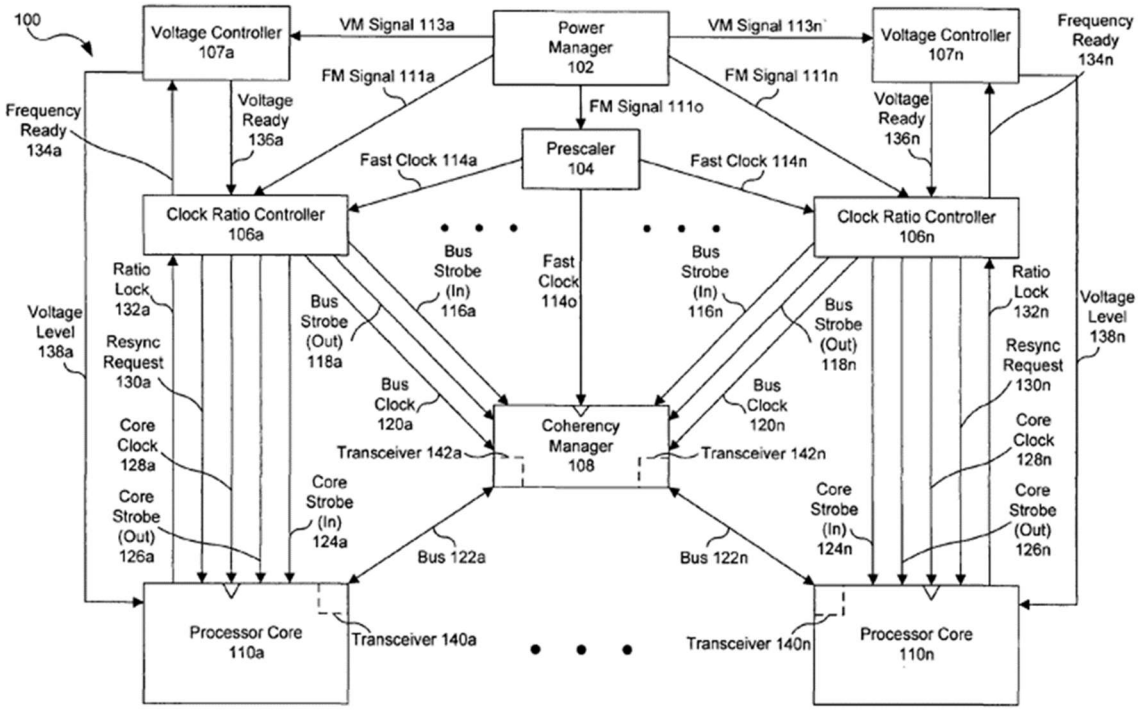


FIG. 1A

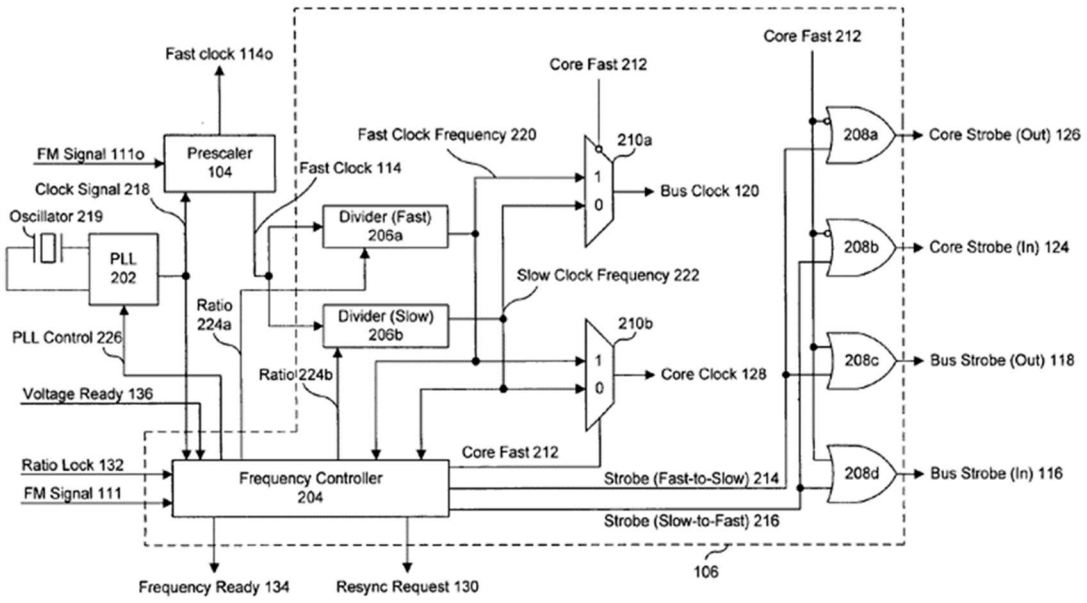


FIG. 2A

*FIGS. 1A and 2A of Knoth (Ex[1005])*

52. In the embodiment of FIG. 1A, digital system 100 is a multiprocessor

digital system that includes at least two processor cores 110. *Id.* at [0023].

> As illustrated in FIG. 1A, in one embodiment, ***digital system 100 is a multiprocessor digital system that includes at least two processor cores 110.*** The present invention, however, is not limited to multiprocessor digital systems. The invention encompasses, for example, any digital system that has a first digital circuit and a second digital circuit coupled together by a bus.

*Id.* at [0023].

53. Power management unit 102 provides individual frequency management (FM) signals 111a-n to clock ratio controllers 106a-n to initiate frequency adjustments and individual management (VM) signals 113a-n to voltage controllers 107a-n to initiate voltage adjustments. *Id.* at [0025].

> In an embodiment, ***power management unit 102 provides individual frequency management (FM) signals 111a-n to clock ratio controllers 106a-n and individual voltage management (VM) signals 113a-n to voltage controllers 107a-n.*** These signals are also referred to herein as power management signals. A FM signal 111o is also provided to prescaler 104. ***These FM signals are used, for example, to initiate frequency adjustments. The VM signals are used, for example, to initiate voltage adjustments.*** FM signals 111a-n can be used to individually control/scale the frequency of each processor core 110a-n. VM signals 113a-n can be used to individually control/scale the voltage of each processor core 110a-n. FM signal 111o can be used to control/scale the frequency of all of the components of digital system 100.

*Id.* at [0025].

54. FM signals 111a-n can be used to individually and independently control/scale the frequency of each processor core 110a-n. VM signals 113a-n can be used to individually and independently control/scale the voltage of each processor

core 110a-n. *Id.*

55. Clock ratio controller 106 in FIG. 2A is coupled to a phase-locked loop (PLL) 202 that outputs a clock signal 218 based on timing pulses generated by an oscillator 219. *Id.* at [0041]-[0042].

> FIGS. 2A-B are diagrams that further illustrate an example clock ratio controller 106 according to an embodiment of the present invention. As shown in FIG. 2A, ***clock ratio controller 106 is coupled to a phase-locked loop (PLL) 202*** and prescaler 104 of digital system 100.
>
> ***PLL 202 outputs a clock signal 218 based on, for example, timing pulses generated by an oscillator 219.*** Clock signal 218 is provided to prescaler 104 and a frequency controller 204. In an embodiment, PLL 202 is controlled by frequency controller 204 using a PLL control signal 226.

*Id.* at [0041]-[0042].

56. Coherency manager 108 and processor cores 110a-n are coupled together by busses 122a-n to facilitate communication between the processor cores. *Id.* at [0006] and [0033].

> In an embodiment, the clock ratio controller outputs strobe signals that enable operation of transmit registers and receive registers of the first digital circuit and the second digital circuit. ***These transmit registers and receive registers are used to facilitate communications between the first digital circuit and the second digital circuit.***

*Id.* at [0006].

> Coherency manager 108 coordinates cache memory access for processor cores 110a-n. As shown in FIG. 1A, ***coherency manager 108 and processor cores 110a-n are coupled together by busses 122a-n.*** In an embodiment, coherency manager 108 runs at the same clock speed as the fastest processor core 110a-n.

*Id.* at [0033].

**B. U.S. Patent No. 8,122,270 B2 to Allarey et al. ("Allarey") (Ex[1006])**

57. U.S. Patent No. 8,122,270 ("Allarey") was filed on September 29, 2008, and issued on February 21, 2012. Its title is "Voltage Stabilization For Clock Signal Frequency Locking." Allarey is therefore prior art to the '339 Patent under at least 35 U.S.C. § 102(e)(2) [Pre-AIA].

58. Allarey discloses stabilizing voltage supplied to a multi-core processor during a clock signal frequency locking process:

> The invention relates to ***stabilizing voltage supplied to a multi-core processor during a clock signal frequency locking process***.

*Id.* at 1:6-8.

> A modern multi-core processor, such as an Intel® architecture processor or another brand processor, generally has multiple power states available to allow for power conservation when the processor is not busy. ***The voltage supplied to the processor and the frequency of the processor may be dynamically modified during operation based on a number of factors such as the current power state of the processor. It is generally beneficial to have a stable and unchanging voltage supplied to a clock signal generation circuit, such as a phase locked loop (PLL), when the PLL is in the process of modifying (e.g. relocking) the frequency of the clock signal being output.*** Asynchronous voltage changes during this time may disrupt a PLL lock process.

*Id.* at 1:12-24.

59. Allarey discloses that the voltage supplied to the processor and the frequency of the processor may be dynamically modified during operation, and that it is beneficial to have a stable voltage supplied to a clock signal generation circuit,

such as a phase lock loop (PLL), when the PLL is in the process of relocking the frequency of the clock signal being output:

> ***The voltage supplied to the processor and the frequency of the processor may be dynamically modified during operation based on a number of factors such as the current power state of the processor. It is generally beneficial to have a stable and unchanging voltage supplied to a clock signal generation circuit, such as a phase locked loop (PLL), when the PLL is in the process of modifying (e.g. relocking) the frequency of the clock signal being output.*** Asynchronous voltage changes during this time may disrupt a PLL lock process.

*Id.* at 1:15-24.

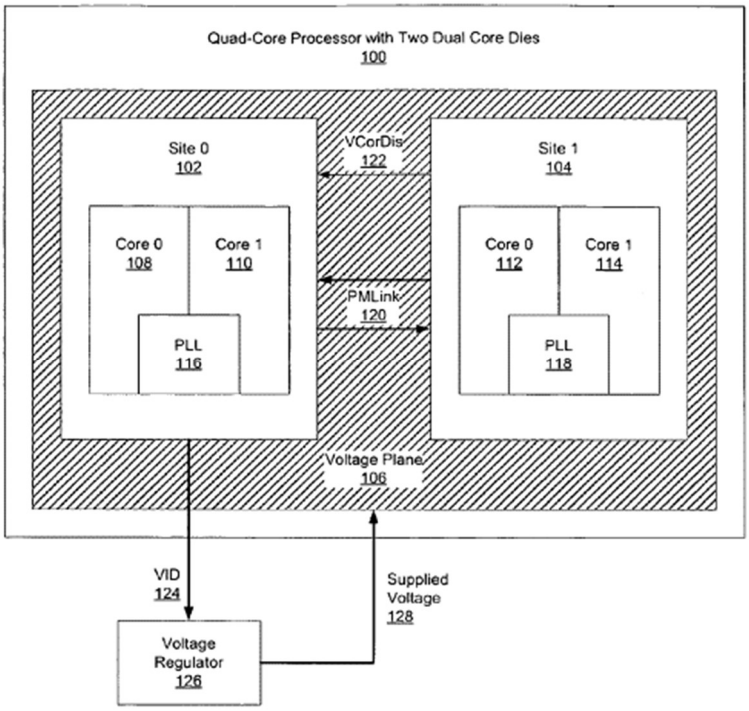


*FIG. 1 of Allarey (Ex[1006])*

60. The quad-core processor 100 of FIG. 1 has two dual-core dies. Site 0 (102) includes two processing cores, core 0 (108) and core 1 (110), and site 1 (104)

also includes two processing cores, core 0 (112) and core 1 (114). *Id.* at 2:41-46.

> ***A quad-core processor with two dual-core dies 100 is portrayed. In many embodiments, the processor 100 includes two sites, site 0 (102) and site 1 (104).*** Both sites are coupled to a common voltage plane 106. ***Site 0 (102) includes two processing cores, core 0 (108) and core 1 (110). Site 1 (104) also includes two processing cores, core 0 (112) and core 1 (114).***

*Id.* at 2:41-46.

### C. U.S. Patent Application Publication No. 2010/0122101 A1 to Naffziger et al. ("Naffziger") (Ex[1010])

61. U.S. Patent Appl. No. 2010/0122101 ("Naffziger") is a published U.S. patent application filed on November 11, 2008, and published on May 13, 2010. Ex[1010] at 1. Its title is "Method and Apparatus for Regulating Power Consumption." Naffziger is therefore prior art to the '339 Patent under at least 35 U.S.C. § 102(e)(1).

62. Naffziger discloses a method for controlling power consumption while maximizing processor performance:

> ***A method for controlling power consumption while maximizing processor performance.*** The method includes, for a time interval of operation in a first operational state, determining an amount of power consumed during by one or more cores of a processor, calculating, a power error based on the amount of power consumed in the time interval, obtaining a power error term for the interval by adding the power error to a power error term from a previous time interval, and comparing the power error term to at least a first error threshold. If the power error term is outside a range defined at least in part by the first error threshold, the method exits the first operational state and enters a second operational state. If the power error term is within the range

> defined at least in part by the first error threshold, operation continues in the first operational state.

*Id.* at Abstract.

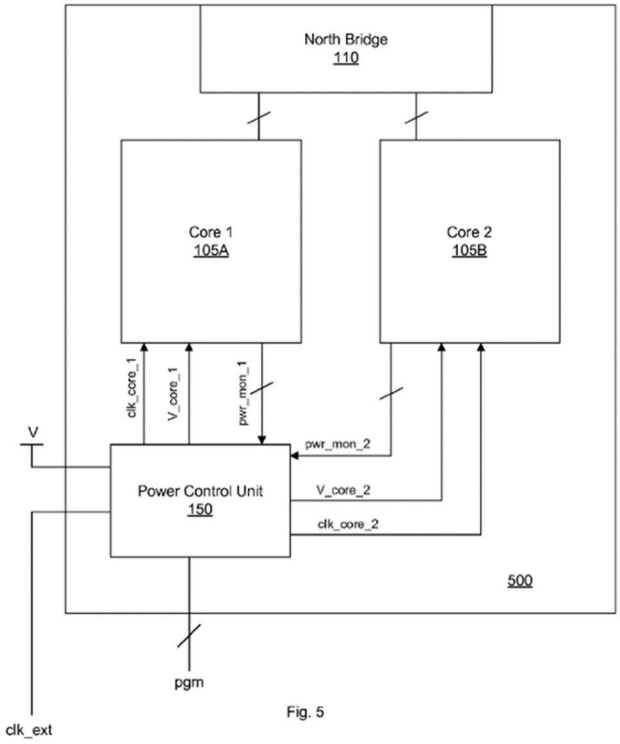


*FIG. 5 of Naffziger (Ex[1010])*

63. In the embodiment of FIG. 5, processor 500 is a dual core processor, having a first core 105A and a second core 105B:

> Turning now to FIG. 5, a block diagram of one embodiment of a processor having a power control unit and a plurality of processor cores is shown. In the embodiment shown, ***processor 500 is a dual core processor, having a first core 105A and a second core 105B.*** Processor 500 also includes a north bridge unit 110 and a power control unit 150, both of which are coupled to both processor cores 105A and 105B. Power control unit 150 is configured to monitor the power of both processor cores 105A and 105B, and configured to alternate operation of these cores between at least a high performance state and a low power state. As noted above, additional, intermediate operational states may be implemented. Power control unit 150 may accomplish the

> switching between operational states of processor cores 105A and 105B by changing the frequency of clock signals provided to the cores, changing supply voltages provided to the cores, or both.

*Id.* at [0053].

64. Changing operational states may include adjusting a supply voltage, changing a clock frequency, or any other parameter that may be changed to affect performance and/or power consumption:

> ***Changing operational states may include adjusting a supply voltage, changing a clock frequency, or any other parameter that may be changed to affect performance and/or power consumption. In one embodiment, entering a high performance state from a low power state includes increasing a core supply voltage and/or a core clock frequency.*** Such parameters may be adjusted by a power control unit that is configured to monitor the power, determine a power error for a given time interval, compute the power error term for the given time interval, and compare the power error term to one or more error thresholds associated with the operational state. The power control unit may also be configured to cause the processor to change operational states based on the comparisons.

*Id.* at [0012].

65. Further, cores 105A and 105B have separate, independent core supply voltage planes and/or are configured to receive separate, independently controlled core clock signals:

> In some embodiments, the operational state of cores 105A and 105B are in conjunction with each other. That is, both processor cores 105A and 105B are operated in the same state. However, embodiments are possible and contemplated wherein ***the operational states of processor cores 105A and 105B may be controlled independently of one another***. For example, in such an embodiment, processor core 105A may operate in the high performance state, while processor core 105B operates in the low power state. ***Such a configuration may require that***

> ***cores 105A and 105B have separate, independent core supply voltage planes and/or are configured to receive separate, independently controlled core clock signals.*** In addition, embodiments are possible and contemplated wherein power control unit 150 is configured to monitor the processing workloads of each processor core 105A and 105B, and may change operational states of the cores by changing their respective processing workloads.

*Id.* at [0054].

> Another embodiment is contemplated wherein ***a multi-core processor includes multiple power control units, each of which is associated with a corresponding one of a plurality of processor cores***. In such an embodiment, ***each power control unit may separately control the states of operation of its corresponding core.*** The power control unit for each core may also monitor power consumption of the processor as a whole, which may be used, along with one or more power error terms, as a basis for determining the operational state of its corresponding processor core.

*Id.* at [0055].

66. Additionally, clock control unit may be implemented using a phase locked loop (PLL) that can be used for adjusting the frequency of a clock signal:

> ***Clock control unit 210 may be implemented using a phase locked loop (PLL)***, a delay locked loop (DLL), or other type of circuitry ***that can be used for adjusting the frequency of a clock signal.*** In the embodiment shown, an external clock signal, clk_ext, is provided at a first frequency to clock control unit 210. A core clock signal, clk_core, is provided by clock control unit 210 to processor core 105. Clock control unit 210 is configured to vary the frequency of the core clock signal based on signals received from control logic 220. In one embodiment, control logic 220 may cause processor 100 to enter the low power state by directing clock control unit 210 to reduce the frequency of the core clock signal, and may cause the processor 100 to enter the high performance state by directing clock control unit to increase the frequency of the core clock signal.

*Id.* at [0033].

67. Power control unit 150 is configured to monitor the power of both processor cores 105A and 105B and configured to alternate operation of these cores between at least a high-performance state and a low power state:

> Turning now to FIG. 5, a block diagram of one embodiment of a processor having a power control unit and a plurality of processor cores is shown. In the embodiment shown, processor 500 is a dual core processor, having a first core 105A and a second core 105B. Processor 500 also includes a north bridge unit 110 and a power control unit 150, both of which are coupled to both processor cores 105A and 105B. ***Power control unit 150 is configured to monitor the power of both processor cores 105A and 105B, and configured to alternate operation of these cores between at least a high performance state and a low power state. As noted above, additional, intermediate operational states may be implemented. Power control unit 150 may accomplish the switching between operational states of processor cores 105A and 105B by changing the frequency of clock signals provided to the cores, changing supply voltages provided to the cores, or both.***

*Id.* at [0053].

**D. U.S. Patent Application Publication No. 2005/0034002 A1 to Flautner ("Flautner") (Ex[1007])**

68. U.S. Patent Appl. No. 2005/0034002 ("Flautner") is a published U.S. patent application filed on August 4, 2003, and published on February 10, 2005. Ex[1007] at 1. Its title is "Performance Control Within a Multi-Processor System." Flautner is therefore prior art to the '339 Patent under at least 35 U.S.C. § 102(b).

69. Flautner discloses a multi-core processor that includes a synchronisation module 50 and a voltage level shifter 52 to deal with clock synchronization issues and the different supply voltage levels between cores in two

independent domains. *Id.* at [0047]-[0048], FIG. 4.

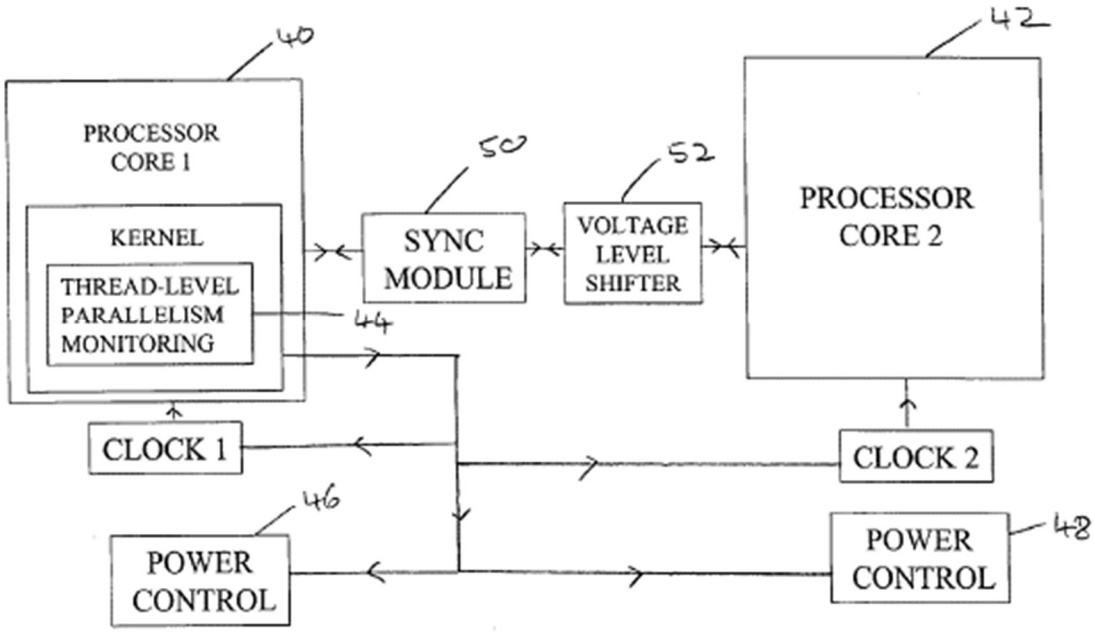


*FIG. 4 of Flautner (Ex[1007])*

70. The embodiment of FIG. 4 provides a first processor core 40 and a second processor core 42 comprising a multi-processing system:

> When using different clock speeds and voltage levels in ***the first processor core 40 and the second processor core 42*** it will be appreciated that a synchronisation module 50 and a voltage level shifter 52 is provided between the first processor core 40 and the second processor core 42 to deal with clock synchronisation issues and the different supply voltage levels (voltage signalling levels) between the two domains.

*Id.* at [0048].

71. A synchronization module 50 and a voltage level shifter 52 is provided between the first processor core 40 and the second processor core 42 to deal with clock synchronization issues and the different supply voltage levels between the two

domains:

> When using different clock speeds and voltage levels in the first processor core 40 and the second processor core 42 it will be appreciated that ***a synchronisation module 50 and a voltage level shifter 52 is provided between the first processor core 40 and the second processor core 42 to deal with clock synchronisation issues and the different supply voltage levels (voltage signalling levels) between the two domains***.

*Id.* at [0048].

**E. U.S. Patent Application Publication No. 2007/0080696 A1 to Kumar et al. ("Kumar") (Ex[1009])**

72. U.S. Patent Appl. No. 2007/0080696 ("Kumar") is a published U.S. patent application filed on October 11, 2005, and published on April 12, 2007. Ex[1009] at 1. Its title is "Integrated Circuit Package Resistance Measurement." Kumar is therefore prior art to the '339 Patent under at least 35 U.S.C. § 102(b).

73. Kumar discloses measuring a resistance of a package of an integrated circuit:

> For one embodiment, an integrated circuit includes a node to couple one or more components to the integrated circuit to carry current through a package for the integrated circuit. ***The integrated circuit also includes a monitor to measure a resistance of the package*** based at least in part on a reference resistance of the package and a resistance of one or more components that are to carry current through the package. For another embodiment, current through one or more components that are to carry current through a package for an integrated circuit is controlled. A resistance of the package is measured based at least in part on a reference resistance of the package and a resistance of one or more components that are to carry current through the package.

*Id.* at Abstract.

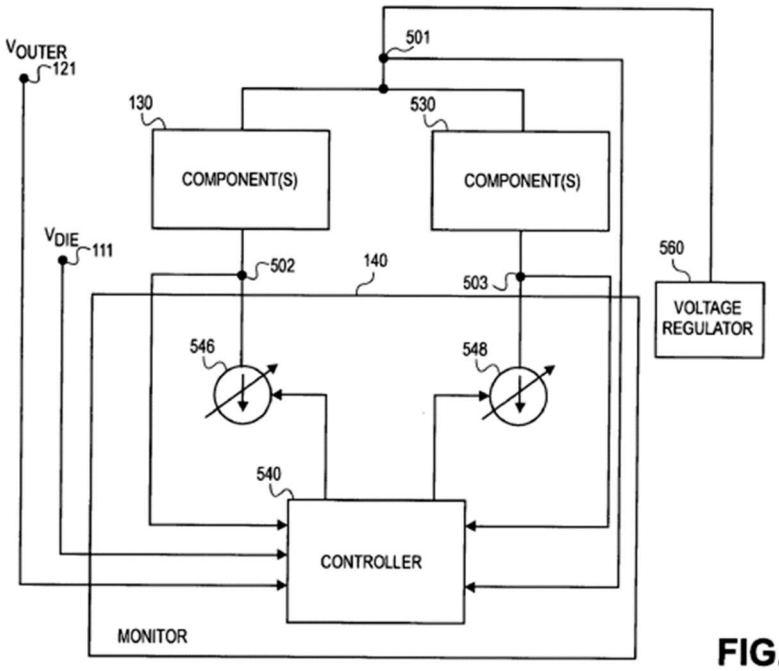


FIG. 5

*FIG. 5 of Kumar (Ex[1010])*

74. The embodiment of FIG. 5 shows a voltage $V_1$ across one or more components 130 and a voltage $V_2$ across one or more components 530:

> For block 706, ***controller 540 may identify whether a voltage V1 across one or more components 130 and a voltage V2 across one or more components 530 satisfy one or more predetermined relationships.*** Controller 540 for one embodiment may identify for block 706 whether the absolute value of the difference between voltage V1 and voltage V2 is less than, or less than or equal to, a predetermined amount or a predetermined percentage of either voltage V1 or voltage V2. Controller 540 for one embodiment may identify for block 706 whether voltage V1 and voltage V2 are substantially equal. For one embodiment where one or more components 130 and one or more components 530 are coupled to a common supply voltage node 501, controller 540 for one embodiment may identify for block 706 whether measured voltages from nodes 502 and 503 satisfy one or more predetermined relationships.

*Id.* at [0058].

75. Kumar teaches that the voltage $V_1$ and the voltage $V_2$ satisfy one or more predetermined relationships, including, for example, whether the absolute value of the difference between voltage $V_1$ and the voltage $V_2$ is less than, or equal to, a predetermined amount or a predetermined percentage of either $V_1$ and the voltage $V_2$:

> For block 706, ***controller 540 may identify whether a voltage V1 across one or more components 130 and a voltage V2 across one or more components 530 satisfy one or more predetermined relationships.*** Controller 540 for one embodiment may identify for block 706 whether the absolute value of the difference between voltage V1 and voltage V2 is less than, or less than or equal to, a predetermined amount or a predetermined percentage of either voltage V1 or voltage V2. Controller 540 for one embodiment may identify for block 706 whether voltage V1 and voltage V2 are substantially equal. For one embodiment where one or more components 130 and one or more components 530 are coupled to a common supply voltage node 501, controller 540 for one embodiment may identify for block 706 whether measured voltages from nodes 502 and 503 satisfy one or more predetermined relationships.

*Id.* at [0058].

> If voltage V1 and voltage V2 do not satisfy one or more predetermined relationships for block 706, controller 540 for block 708 may control programmable current source 546 to adjust current I1 through one or more components 130 and/or control programmable current source 548 to adjust current I2 through one or more components 530. ***Controller 540 for one embodiment for block 708 may adjust current I1 and/or current I2 in any suitable manner to help voltage V1 and voltage V2 satisfy one or more predetermined relationships for block 706. Controller 540 may repeat operations for blocks 706 and 708 until voltage V1 and voltage V2 satisfy one or more predetermined relationships for block 706.***

*Id.* at [0059].

### F. U.S. Patent Application Publication No. 2011/0153984 A1 to Wolfe et al. ("Wolfe") (Ex[1008])

76. U.S. Patent Appl. No. 2011/0153984 ("Wolfe") is a U.S. patent application filed by Empire Technology Development LLC on December 21, 2009, and published on June 23, 2011. Ex[1008] at 1. Wolfe is therefore prior art to the '339 Patent under at least 35 U.S.C. § 102(e)(1).

77. Wolfe teaches a grid organization of processor cores—the processor cores of a multi-core processor may be "arranged in rows and columns in a 2-dimensional array":

> FIG. 1A is an illustration of a processor core 100 in a multi-core processor 102 adapted to operate based on a selected voltage level and a clock frequency, arranged in accordance to at least some embodiments of the present disclosure. ***The multi-core processor 100 may include multiple processor cores arranged in rows and columns in a 2-dimensional array.*** The processor cores may be interconnected by a communication network using switches, which is illustrated in FIG. 3. A processor core may be supplied a certain voltage level (e.g., a selected voltage level 124) from a voltage control circuit (e.g., a voltage control circuit 110) and a clock signal at a certain clock frequency (e.g., a selected clock frequency 126) from a clock control circuit (e.g., a clock control circuit 112).

*Id.* at [0014].

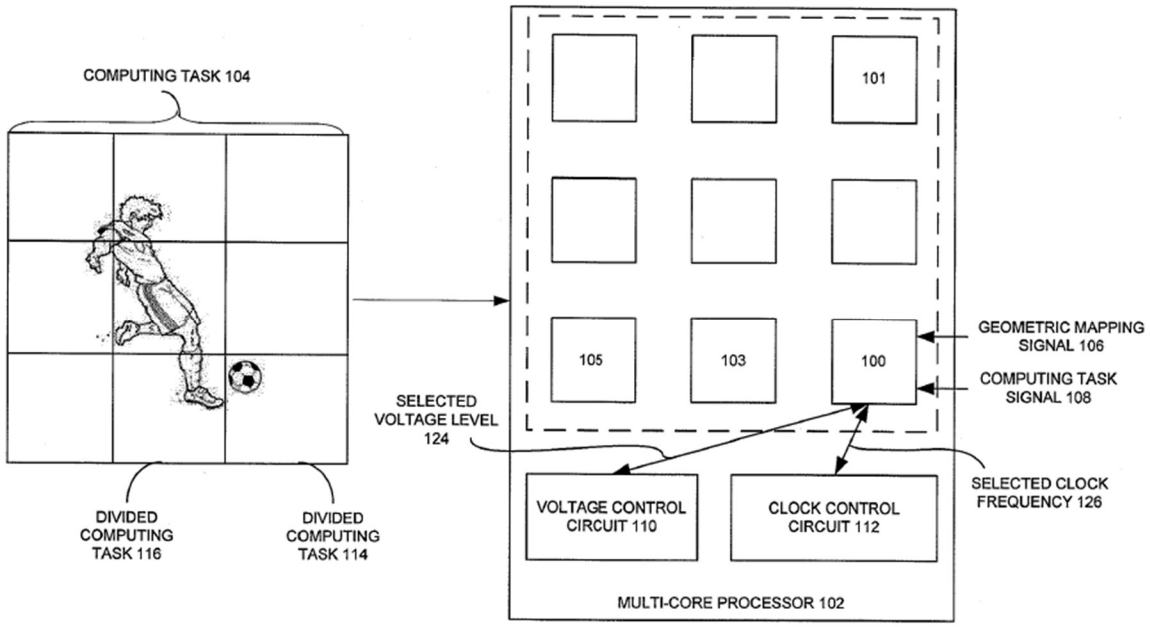


*FIGURE 1A of Wolfe (Ex[1008])*

78. Wolfe also discloses grouping cores based on geometric mapping:

> When the computing task 104 is divided, a geometric mapping signal 106 and a computing task signal 108 may be generated by a host processor, or one or more designated processor cores. The computing task signal 108 may convey information associated with one or more of the divided computing tasks, and the geometric mapping signal 106 may convey information associated with the geometric mapping between the divided computing task and a processor core. The information associated with the computing tasks may include pixel data, reference frame data, motion estimation data, motion compensation data, or other data that forms the basis of the computing task. ***The information associated with the geometric mapping may include the assignment of a particular divided computing task to one or more processor cores of the multi-core processor.***

*Id.* at [0016].

> Processing for the method 200 may begin at block 202, ***"Determine First Workload for First Computing Task Geometrically Mapped to First Processor Core,"*** where the method 200 may be arranged to determine the first workload of the first processor core in the multi-core processor. ***The first workload is for the first processor core to perform***

> ***a first computing task associated with a first image area and a first geometric mapping between the first computing task and the first processor core.***

*Id.* at [0022].

## VI. GROUND 1: CLAIMS 1, 5, 8-10, 14 AND 21 ARE OBVIOUS UNDER §103 OVER KNOTH IN VIEW OF ALLAREY

79. Knoth and Allarey, alone or in combination, teach each and every limitation of Claims 1, 5, 8-10, 14 and 21 and therefore render these claims obvious.

### A. Reasons to Combine Knoth and Allarey

80. A PHOSITA would have been motivated to combine Knoth and Allarey and would have had a reasonable expectation of success in doing so, because they relate to the same technology field. Knoth states:

> The present invention generally relates to digital systems. More particularly, ***it relates to a clock ratio controller for dynamic voltage and frequency scaled digital systems***, and applications thereof.

Ex[1005] at [0001].

81. Similarly, Allarey "relates to stabilizing voltage supplied to a multi-core processor during a clock signal frequency locking process." Ex[1006] at 1:6-8. As a result, both Knoth and Allarey are related to digital processor systems for controlling voltage and clock signal frequency.

82. Further, Knoth and Allarey both are drawn to the field of multi-core processors. For example, Knoth has a multi-core processing circuit and multiple processor cores:

> As illustrated in FIG. 1A, in one embodiment, ***digital system 100 is a multiprocessor digital system that includes at least two processor cores 110.*** The present invention, however, is not limited to multiprocessor digital systems. The invention encompasses, for example, any digital system that has a first digital circuit and a second digital circuit coupled together by a bus.

Ex[1005] at [0023].

> FIG. 8 is a diagram of an example digital circuit of system 800 according to an embodiment of the present invention. ***System 800 includes a multi-core processing circuit 802***, a memory 804, an input/output (I/O) controller 806, a clock 808, and custom hardware 810. In an embodiment, system 800 is a system on a chip (SOC) in an application specific integrated circuit (ASIC).

*Id.* at [0083].

83. Similarly, Allarey has a multi-core processor, as discussed below.

> The invention relates to stabilizing voltage supplied to ***a multi-core processor*** during a clock signal frequency locking process.

Ex[1006] at 1:6-8.

84. Furthermore, Knoth and Allarey each address similar problems and propose similar solutions for managing voltage and frequency scaling of multi-core processors, similar to that described in the '339 Patent. For example, Knoth discloses solutions for multi-core processor systems to control power consumption that include changing operational states with different power consumptions by dynamically adjusting the voltage and clock frequency:

> The present invention generally relates to digital systems. More particularly, it relates to ***a clock ratio controller for dynamic voltage and frequency scaled digital systems***, and applications thereof.

> ***Dynamic voltage and frequency scaling are techniques for achieving low power consumption in digital systems while maintaining just enough processing speed to ensure that processing tasks are timely satisfied.*** While these techniques are useful for reducing power consumption, dynamically changing clocking frequencies in conventional digital systems result in loss of synchronization between various system elements. As a result, communications and the exchange of data between various components are disrupted until a resynchronization occurs. What are needed are improved techniques for implementing dynamic voltage and frequency scaling that overcome the limitations noted above.
>
> The present invention provides a clock ratio controller for dynamic voltage and frequency scaled digital systems, and applications thereof. In an embodiment, ***the clock ratio controller is used, for example, to adjust the frequency of clock signals of a digital system without a loss of synchronization between system elements.***

Ex[1005] at [0001]-[0003].

85. Like Knoth, Allarey describes solutions to optimize the ability of multi-core processor systems to conserve power by dynamically modifying the voltage supplied to and the frequency of the processor:

> The invention relates to stabilizing voltage supplied to a multi-core processor during a clock signal frequency locking process.
>
> A modern multi-core processor, such as an Intel® architecture processor or another brand processor, generally has multiple power states available to allow for power conservation when the processor is not busy. ***The voltage supplied to the processor and the frequency of the processor may be dynamically modified during operation based on a number of factors such as the current power state of the processor.*** It is generally beneficial to have a stable and unchanging voltage supplied to a clock signal generation circuit, such as a phase locked loop (PLL), when the PLL is in the process of modifying (e.g. relocking) the frequency of the clock signal being output. Asynchronous voltage changes during this time may disrupt a PLL lock process.

Ex[1006] at 1:6-24.

86. Both of these references are aligned with the '339 Patent, which applies "dynamic supply voltage and clock speed control" for power consumption management. Ex[1001] at 1:6-14.

87. According to the face of the Knoth patent, the inventor of the patent is Matthias Knoth, and the assignee is MIPS Technologies, Inc. of Mountain View, CA. MIPS Technologies, Inc., now known as MIPS Tech LLC ("MIPS"), was a semiconductor design company.

88. According to the face of the Allarey patent, the inventors of the patent are Jose Allarey, Sanjeev Jahagirdar, and Ivan Herrera. The assignee of the patent is Intel Corporation ("Intel") of Santa Clara, CA, a large semiconductor chip manufacturer.

89. A PHOSITA would have looked to publications by such leading companies in semiconductor design, such as MIPS and Intel, and considered the similar techniques disclosed in these references for optimizing the power and/or performance of cores. The similar technologies of MIPS and Intel in the field of semiconductor design could be used to improve power consumption of the multi-core processors.

90. A PHOSITA also would have found the combination obvious to try because it combines well-known techniques that are related—e.g., Knoth provides a

system that can dynamically adjust clock signals without losing synchronization, focusing on coordination between different clock domains; Allarey provides a method to ensure stable voltage during these adjustments, which is crucial for the proper functioning of Knoth's clock ratio controller. As a result, Knoth and Allarey work together to improve functionality of the multi-core processors.

91. Additionally, Knoth's method relies on stable voltage conditions during frequency adjustments, which Allarey's invention directly addresses. For example, in Knoth, "if VM signals 113a-n indicate a voltage decrease that would require a frequency decrease, voltage controllers 107a-n may wait for frequency ready signals 134a-n before adjusting the voltage":

> In an embodiment, ***if VM signals 113a-n indicate a voltage decrease that would require a frequency decrease, voltage controllers 107a-n may wait for frequency ready signals 134a-n before adjusting the voltage.*** If FM signals 111a-n indicate a frequency increase that would require a voltage increase, clock ratio controllers 106a-n may wait for voltage ready signals 136a-n before adjusting the frequency.

Ex[1005] at [0032].

92. Voltage stability is essential during frequency adjustment; as a result, voltage adjustment is applied only after the frequency is stable, as represented by the frequency ready signals. Similarly, Allarey states that "[d]uring the PLL frequency locking window of time, it is greatly beneficial that the voltage supplied to the site the PLL feedback loop circuitry is located within remains stable":

> In many embodiments, due to power conservation logic within processor 100 such as Enhanced Intel® SpeedStep® Technology or other processor power management technology, each site with processor 100 might actively modify the frequency of the cores if the cores are switching between a sleep mode, a low frequency mode, a high frequency mode, or another such frequency-changing mode. ***In these embodiments, at any given time, processor power management logic may need to modify the frequency of the clock signal being supplied to the cores by PLL 116 and PLL 118. A PLL frequency locking process is not instantaneous and instead requires a finite window of time.*** The PLL locking (or re-locking) process requires a feedback loop circuit to help modify the PLL frequency. ***The feedback loop circuitry for each PLL is affected by a core voltage change, which results in longer locking (re-locking) times. During the PLL frequency locking window of time, it is greatly beneficial that the voltage supplied to the site the PLL feedback loop circuitry is located within remains stable. A PLL frequency locking process potentially will not succeed or take a longer period of time if there is a change in the voltage supplied to the PLL feedback loop circuitry.*** Thus, the PLLs in both site 0 (102) and site 1 (104) benefit if they are given a window of time in which they can be sure that the voltage they are supplied will not change.

Ex[1006] at 3:25-49.

93. A PHOSITA would have found synergy in combining dynamic power adjustments from Knoth with voltage stabilization from Allarey to improve overall processor performance and stability. As discussed above, both Knoth and Allarey discuss the benefits of voltage stabilization during frequency switching.

94. Finally, a PHOSITA would have considerable expectation of success when combining these teachings because the combination would amount to a mere substitution of one known element for another, applying a known technique to a known system ready for improvement, and/or use of known techniques to improve

similar systems.

> In many embodiments, due to power conservation logic within processor 100 such as Enhanced Intel® SpeedStep® Technology or other processor power management technology, each site with processor 100 might actively modify the frequency of the cores if the cores are switching between a sleep mode, a low frequency mode, a high frequency mode, or another such frequency-changing mode. ***In these embodiments, at any given time, processor power management logic may need to modify the frequency of the clock signal being supplied to the cores by PLL 116 and PLL 118. A PLL frequency locking process is not instantaneous and instead requires a finite window of time.*** The PLL locking (or re-locking) process requires a feedback loop circuit to help modify the PLL frequency. ***The feedback loop circuitry for each PLL is affected by a core voltage change, which results in longer locking (re-locking) times. During the PLL frequency locking window of time, it is greatly beneficial that the voltage supplied to the site the PLL feedback loop circuitry is located within remains stable. A PLL frequency locking process potentially will not succeed or take a longer period of time if there is a change in the voltage supplied to the PLL feedback loop circuitry.*** Thus, the PLLs in both site 0 (102) and site 1 (104) benefit if they are given a window of time in which they can be sure that the voltage they are supplied will not change.

Ex[1006] at 3:25-49.

95. As a result, having a constant voltage allows the frequency locking process to successfully and more quickly stabilize. Therefore, the teachings and considerations of Allarey would allow a PHOSITA to improve on Knoth's systems effortlessly (and vice versa). For at least these reasons, a PHOSITA would have been motivated to combine Knoth and Allarey.

## B. Independent Claim 1

### 1. 1[pre] - A multi-core processor, comprising:

96. To the extent that the preamble is limiting, Knoth teaches this subject matter. For example, Knoth discloses that "digital system 100 is a multiprocessor digital system that includes at least two processor cores 110":

> As illustrated in FIG. 1A, in one embodiment, ***digital system 100 is a multiprocessor digital system that includes at least two processor cores 110.*** The present invention, however, is not limited to multiprocessor digital systems. The invention encompasses, for example, any digital system that has a first digital circuit and a second digital circuit coupled together by a bus.

Ex[1005] at [0023].

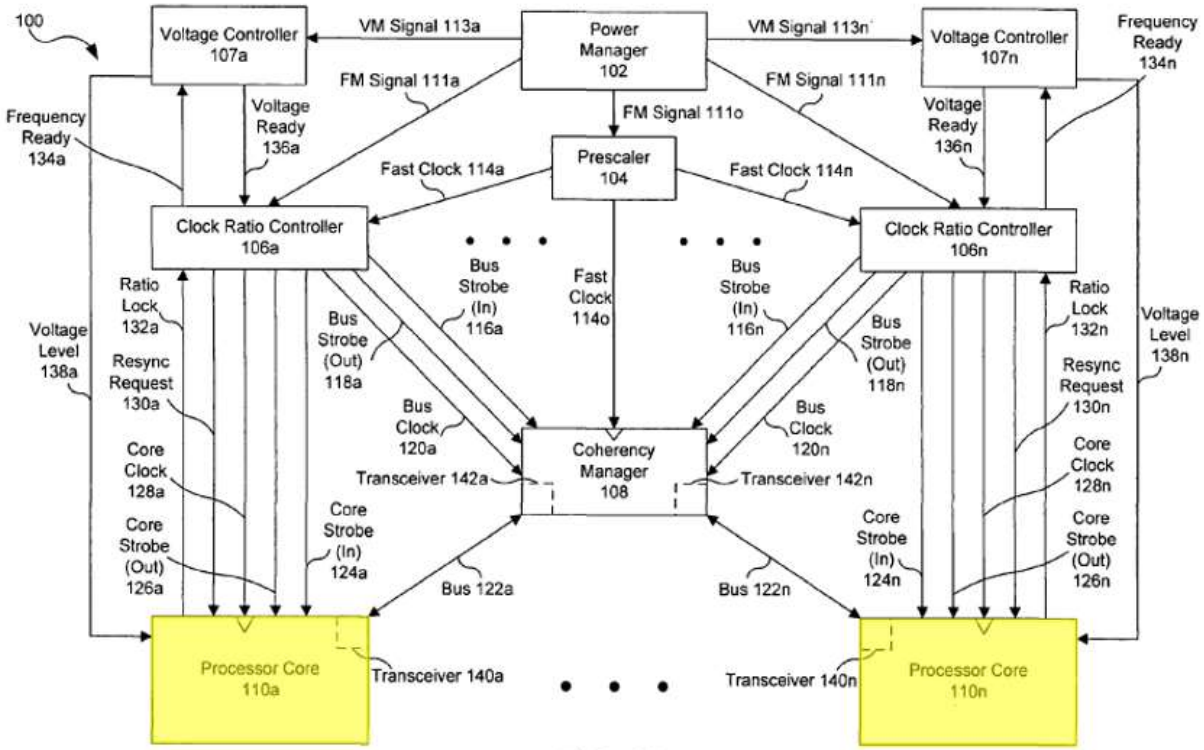


FIG. 1A

*FIG. 1A of Knoth, annotated*

97. The multiprocessor digital system 100 includes at least two processor cores 110a-n (highlighted in yellow) in a disclosed arrangement in which there could be an arbitrary number of cores a through n. A multiprocessor having at least two processor cores discloses the recited "multi-core processor." A multi-core processor is defined as a microprocessor with two or more separate central processing units, or processor cores. Each processor core 110a-n is a separate central processing unit. With the microprocessor 100 having at least two processor cores 110a-n, Knoth teaches the recited "multi-core processor." Knoth also discloses a "multi-core processing circuit":

> ***Multi-core processing circuit 802 is any multiple core processing circuit that includes features of the present invention described herein and/or that implements a method embodiment of the present invention.*** In one embodiment, each processor core of multi-core processing circuit 802 includes an instruction fetch unit, an instruction cache, an instruction decode and dispatch unit, one or more instruction execution unit(s), a data cache, a register file, and a bus interface unit.

*Id.* at [0084].

98. A multi-core processing circuit is a circuit that is used for a multi-core processor. Therefore, Knoth discloses the multi-core processor as recited in the preamble of Claim 1.

99. Allarey also teaches the subject matter of the preamble. Allarey discloses a "multi-core processor":

> The invention relates to stabilizing voltage supplied to ***a multi-core processor*** during a clock signal frequency locking process.

Ex[1006] at 1:6-8.

100. For example, Allarey discloses a quad-core processor 100 that includes two dual-core dies:

> FIG. 1 is an illustration of an apparatus to stabilize a supplied voltage during a clock signal frequency locking process according to some embodiments.
>
> ***A quad-core processor with two dual-core dies 100 is portrayed.*** In many embodiments, ***the processor 100 includes two sites, site 0 (102) and site 1 (104).*** Both sites are coupled to a common voltage plane 106. ***Site 0 (102) includes two processing cores, core 0 (108) and core 1 (110). Site 1 (104) also includes two processing cores, core 0 (112) and core 1 (114).*** Each core includes logic to execute instructions. ***Thus, combined, the two sites have a total of four cores, hence the processor being a quad-core processor with two dual-core dies.*** Each site includes a phase locked loop (PLL) clock signal generation circuit, PLL 116 for site 0 (102) and PLL 118 for site 1 (104). Each PLL is capable of generating a clock signal that the cores located at each respective site can use as a reference clock. Additionally, each PLL can change the frequency of the clock signal through a relocking process. In other embodiments that are not shown, an alternative form of clock signal generation logic generates the clock signal supplied to the cores within the processor 100.

*Id.* at 2:38-58.

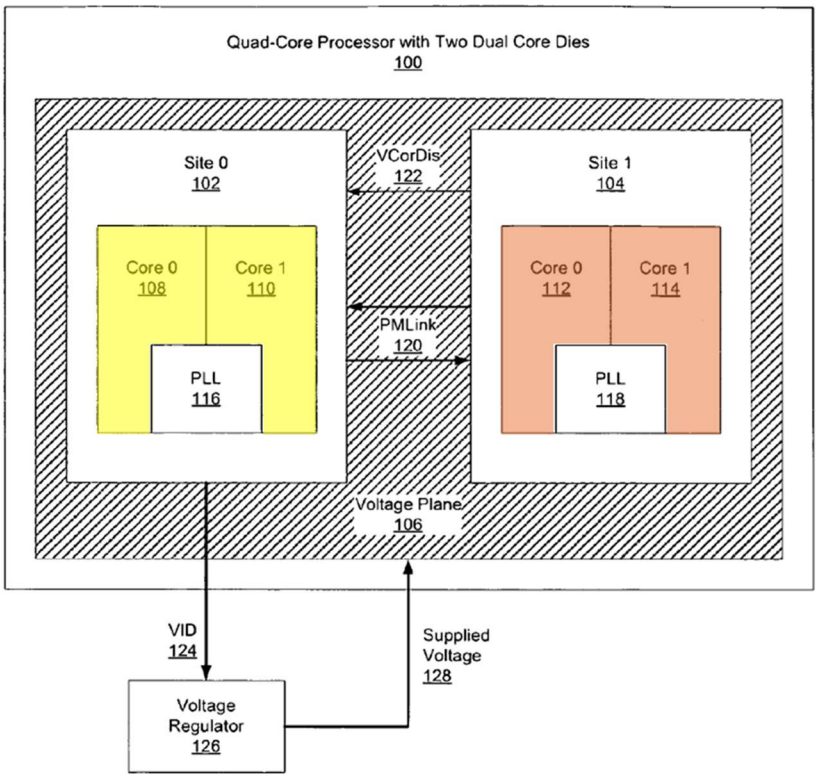


*FIG. 1 of Allarey, annotated*

101. The first dual-core die in site 0 (102) includes a dual-core die comprising core 0 (108) and core 1 (110) (highlighted in yellow) and the second dual-core die in site 1 (104) includes a dual-core die comprising core 0 (112) and core 1 (114) (highlighted in orange), as depicted below in annotated FIG. 1 of Allarey. In Allarey, the processor 100 has core 0 (108), core 1 (110), core 0 (112), and core 1 (114). Accordingly, Allarey teaches a multi-core processor.

102. As another example, Allarey discloses an 8-core processor with two 4-core dies 100 (as depicted below, highlighted in yellow and orange, respectively). Accordingly, Allarey teaches a multi-core processor.

> ***An 8-core processor with two 4-core dies 100 is portrayed. In many embodiments, the processor 200 includes two sites, site 0 (302) and site 1 (304). Both sites are coupled to a common voltage plane 306. Site 0 (302) includes four processor cores, core 0 (308), core 1 (310), core 2 (312), and core 3 (314). Site 1 (304) also includes four processor cores, core 4 (316), core 5 (318), core 6 (320), and core 7 (322).*** Each core includes logic to execute instructions. Thus, combined, the two sites have a total of 8 cores, hence the processor being an 8-core processor with two 4-core dies. Each site includes a PLL feedback loop circuit, PLL 324 for site 0 (302) and PLL 326 for site 1 (304), to help generate a clock signal. Each PLL is capable of generating a clock signal that the core located at each respective site can use as a reference clock. Additionally, each PLL can change the frequency of the clock signal through a re-locking process.

*Id.* at 5:17-33.

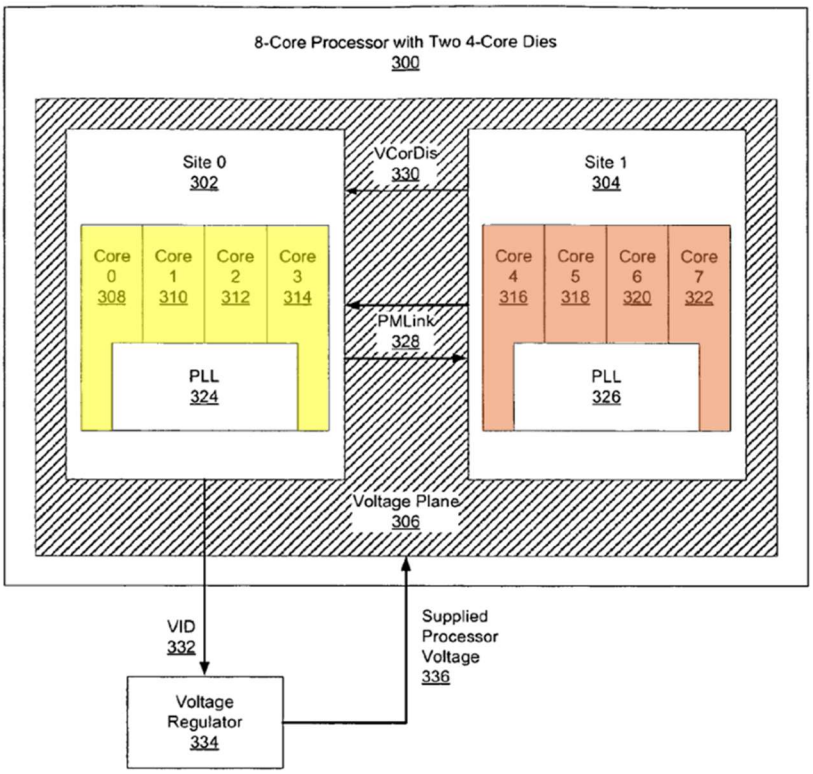


*FIG. 3 of Allarey, annotated*

### 2. 1[a1] - a first set of processor cores of the multi-core processor,

103. Knoth and Allarey, alone or in combination, disclose "a first set of processor cores of the multi-core processor" and "a second set of processor cores of the multi-core processor." A set of processor cores of the multi-core processor comprises one or more processor cores. A set of processor cores can have one processor core, or alternatively, it could have two or more processor cores.

104. Knoth discloses that "digital system 100 is a multiprocessor digital system that includes at least two processor cores 110":

> As illustrated in FIG. 1A, in one embodiment, ***digital system 100 is a multiprocessor digital system that includes at least two processor cores 110.*** The present invention, however, is not limited to multiprocessor digital systems. The invention encompasses, for example, any digital system that has a first digital circuit and a second digital circuit coupled together by a bus.

Ex[1005] at [0023].

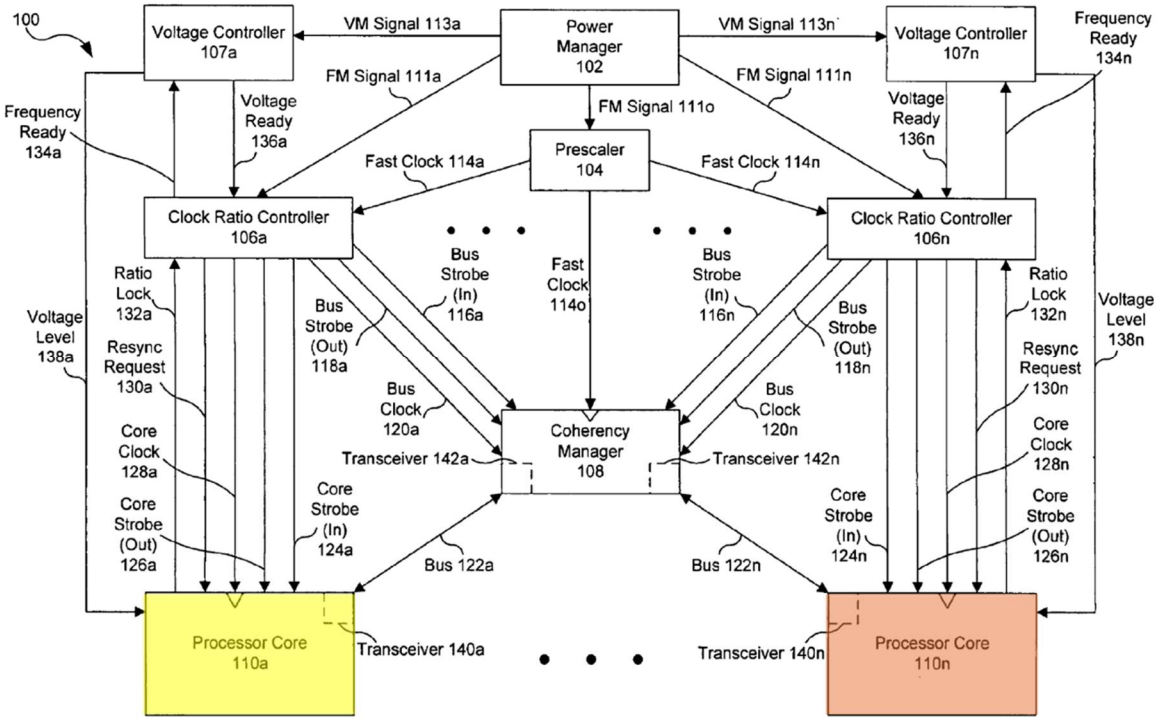


*FIG. 1A of Knoth, annotated*

105. The multiprocessor digital system 100 includes a series of processor cores a-n, with an explicit depiction of at least two processor cores 110a (highlighted in yellow) and 110n (highlighted in orange). Knoth discloses that the processor cores 110a-n each operate independently through dedicated power and frequency management:

> In an embodiment, ***power management unit 102 provides individual frequency management (FM) signals 111a-n to clock ratio controllers 106a-n and individual voltage management (VM) signals 113a-n to voltage controllers 107a-n.*** These signals are also referred to herein as power management signals. A FM signal 111o is also provided to prescaler 104. ***These FM signals are used, for example, to initiate frequency adjustments. The VM signals are used, for example, to initiate voltage adjustments. FM signals 111a-n can be used to individually control/scale the frequency of each processor core 110a-n. VM signals 113a-n can be used to individually control/scale the***

***voltage of each processor core 110a-n.*** FM signal 111o can be used to control/scale the frequency of all of the components of digital system 100.

Prescaler 104 is used to determine the fastest clocking signal that is used in digital system 100. Prescaler 104 provides this clock signal to clock ratio controllers 106a-n and coherency manager 108 via fast clock signals 114a-o. In an embodiment, fast clock signal 114o is equal in frequency to the clocking signal that is used to clock the processor core 110 running the fastest. In an embodiment, when an FM signal 111 associated with the fastest running processor core 110 is adjusted, power manager 102 sends FM signal 111o to prescaler 104 to notify prescaler 104 of the change. Prescaler 104 adjusts fast clock signals 114a-o accordingly.

***Clock ratio controllers 106a-n generate the clock or clocking signals used to control processor cores 110a-n*** and data communications on busses 122a-n. As shown in FIG. 1A, in an embodiment, each clock ratio controller 106a-n generates a core clock signal 128 used to drive an associated processor core 110 and a bus clock signal 120 used to drive an associated transmitter/receiver or transceiver 142 of coherency manager 108. These clocking signals are based on the fast clock signal 114a-n received by clock ratio controllers 106a-n. Core clock signals 128a-n and bus clock signals 120a-n are not required to operate at the same frequency, nor do core clock signals 128a-n have to be faster or slower than bus clock signals 120a-n.

*Id.* at [0025]-[0027].

***Voltage controllers 107a-n control the voltage levels used to run processor cores 110a-n.*** As shown in FIG. 1A, in an embodiment, each voltage controller 107a-n sends a voltage ready signal 136a-n to each clock ratio controller 106a-n to acknowledge when a new voltage level has been set. Each voltage controller 107a-n also sends a voltage or voltage level signal 138a-n to each processor core 110a-n to power each processor core 110a-n.

*Id.* at [0031].

106. The processor cores 110a and 110n disclose the recited first set of

processor cores and the second set of processor cores of the multi-core processor, to the extent that a "set" can include one or more. The processor core 110a is the recited first set of processor cores, and the processor core 110n is the recited second set of processor cores.

107. To the extent it is argued that the processor cores 110a-n are individual cores, or that a "set" means "two or more," it would have been obvious to a PHOSITA to extend the disclosed architecture of Knoth to include two distinct sets of cores, where each set of cores is managed separately. In semiconductor design, extending from one processor core to two processor cores is well-known in the art. Each processor core of the set would receive the same voltage and clock signals. After all, as discussed above, Knoth discloses an arbitrary number of processor cores 1-n and, therefore, more than two cores could be included in two distinct sets. For example, instead of managing individual cores, the system could treat processor core 110a and additional cores associated with it as a first set (e.g., cores b-i), while processor core 110n and its associated cores (e.g., cores j-n) could form a second set. Each set would have been managed independently by the existing dedicated power and frequency management components from Knoth, allowing the system to fine-tune power consumption based on the specific demands of each set. Grouping processor cores into different sets is a simple modification to Knoth. Extending Knoth's architecture to treat two cores or more in at least two distinct sets would

have been a routine and desired modification for a PHOSITA, requiring only minor adjustments to the existing power and frequency management mechanisms, such as sending the power control signals to an individual core to a shared power domain or voltage plane for a set of cores that operates under the same conditions, which was common and well-understood adjustments a PHOSITA would have easily implemented. *Id.* A PHOSITA would have also understood that extending from individual cores to sets of cores would improve the system by enabling more efficient management of power and resources through grouped control, reducing complexity and enhancing scalability for handling diverse workloads. Accordingly, Knoth discloses and/or renders obvious "a first set of processor cores of the multi-core processor" and "a second set of processor cores of the multi-core processor."

108. Furthermore, a PHOSITA would have been motivated to modify Knoth to have first and second sets of cores, rather than individual first and second cores, based on the teachings of Allarey.

109. Specifically, as reviewed above for Claim 1[Preamble] in Ground 1, Allarey discloses a first set of processor cores in the form of a multi-core die in site 0 (highlighted in yellow) and a second set of processor cores in the form of a multi-core die in site 1 (highlighted in orange), as depicted below in FIGS. 1 and 3:

> ***A quad-core processor with two dual-core dies 100 is portrayed. In many embodiments, the processor 100 includes two sites, site 0 (102) and site 1 (104). Both sites are coupled to a common voltage plane 106. Site 0 (102) includes two processing cores, core 0 (108) and core***

> ***1 (110). Site 1 (104) also includes two processing cores, core 0 (112) and core 1 (114).*** Each core includes logic to execute instructions. Thus, combined, the two sites have a total of four cores, hence the processor being a quad-core processor with two dual-core dies. Each site includes a phase locked loop (PLL) clock signal generation circuit, PLL 116 for site 0 (102) and PLL 118 for site 1 (104). Each PLL is capable of generating a clock signal that the cores located at each respective site can use as a reference clock. Additionally, each PLL can change the frequency of the clock signal through a relocking process. In other embodiments that are not shown, an alternative form of clock signal generation logic generates the clock signal supplied to the cores within the processor 100.

Ex[1006] at 2:41-58.

> ***An 8-core processor with two 4-core dies 100 is portrayed. In many embodiments, the processor 200 includes two sites, site 0 (302) and site 1 (304). Both sites are coupled to a common voltage plane 306. Site 0 (302) includes four processor cores, core 0 (308), core 1 (310), core 2 (312), and core 3 (314). Site 1 (304) also includes four processor cores, core 4 (316), core 5 (318), core 6 (320), and core 7 (322).*** Each core includes logic to execute instructions. Thus, combined, the two sites have a total of 8 cores, hence the processor being an 8-core processor with two 4-core dies. Each site includes a PLL feedback loop circuit, PLL 324 for site 0 (302) and PLL 326 for site 1 (304), to help generate a clock signal. Each PLL is capable of generating a clock signal that the core located at each respective site can use as a reference clock. Additionally, each PLL can change the frequency of the clock signal through a re-locking process.

*Id.* at 5:17-33.

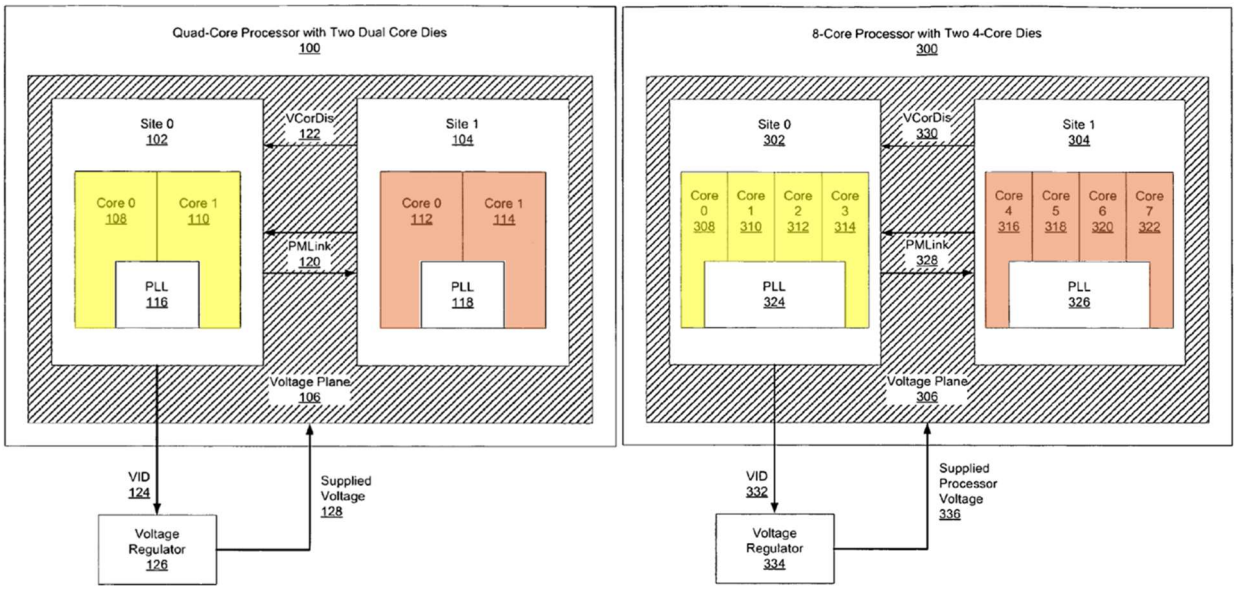


*FIGs. 1 and 3 of Allarey, annotated*

**3. 1[a2] - wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage and a first output clock signal of a first phase lock loop (PLL) having a first clock signal as input;**

110. Knoth teaches this claim limitation. As depicted in FIG. 1A below, each of the processor cores 110a-n receives its respective voltage supply "voltage level signal 138a-n" (shown by the blue line) from its respective voltage controller 107a-n and its respective clock signal core clock signal 128a-n (shown by the purple line) output from respective clock ratio controller 106a-n (shown by the purple box):

> ***Clock ratio controllers 106a-n generate the clock or clocking signals used to control processor cores 110a-n*** and data communications on busses 122a-n. As shown in FIG. 1A, in an embodiment, each clock ratio controller 106a-n generates a core clock signal 128 used to drive an associated processor core 110 and a bus clock signal 120 used to drive an associated transmitter/receiver or transceiver 142 of coherency manager 108. These clocking signals are based on the fast clock signal 114a-n received by clock ratio controllers 106a-n. Core clock signals

128a-n and bus clock signals 120a-n are not required to operate at the same frequency, nor do core clock signals 128a-n have to be faster or slower than bus clock signals 120a-n.

Ex[1005] at [0027].

> ***Voltage controllers 107a-n control the voltage levels used to run processor cores 110a-n.*** As shown in FIG. 1A, in an embodiment, each voltage controller 107a-n sends a voltage ready signal 136a-n to each clock ratio controller 106a-n to acknowledge when a new voltage level has been set. Each voltage controller 107a-n also sends a voltage or voltage level signal 138a-n to each processor core 110a-n to power each processor core 110a-n.

*Id.* at [0031].

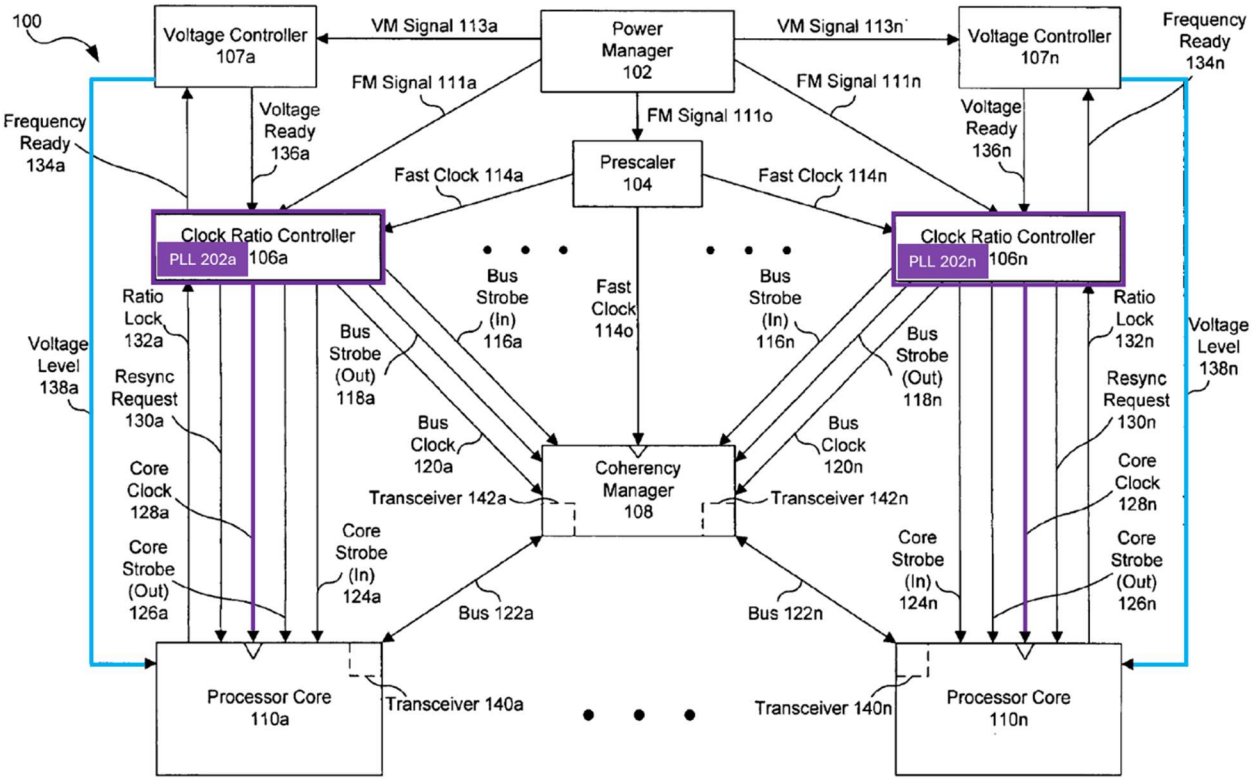


*FIG. 1A of Knoth, modified to show PLL 202a and PLL 202n within clock ratio controller 106a and clock ratio controller 106n*

111. Each of clock ratio controllers 106a-n includes a phase locked loop (PLL) (for ease of discussion, referred to as PLL 202a-n corresponding to clock ratio

controller 106a-n) (highlighted in purple). As illustrated in FIG. 2A, below, Knoth shows a representative clock ratio controller 106 that includes a PLL 202 and an oscillator 219. As a result, clock ratio controllers 106a-n each includes a phase locked loop (PLL) 202 (highlighted in purple) having as input "timing pulses generated by an oscillator 219" within the corresponding clock ratio controller 106 (for ease of discussion, referred to as oscillator 219a-n corresponding to clock ratio controller 106a-n). Ex[1005] at [0041]-[0042]. Further details are provided below:

> FIGS. 2A-B are diagrams that further illustrate an example clock ratio controller 106 according to an embodiment of the present invention. As shown in FIG. 2A, ***clock ratio controller 106 is coupled to a phase-locked loop (PLL) 202*** and prescaler 104 of digital system 100.
>
> ***PLL 202 outputs a clock signal 218 based on, for example, timing pulses generated by an oscillator 219.*** Clock signal 218 is provided to prescaler 104 and a frequency controller 204. In an embodiment, PLL 202 is controlled by frequency controller 204 using a PLL control signal 226.

*Id.* at [0041]-[0042].

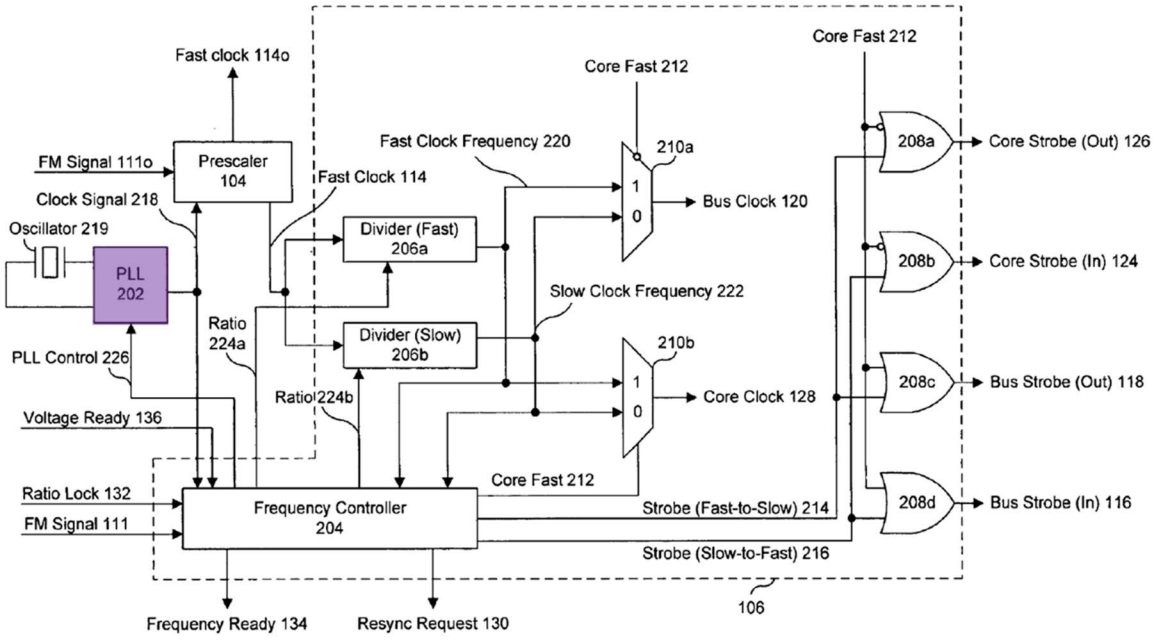


*FIG. 2A of Knoth, annotated*

112. Alternatively, each PLL within clock ratio controller 106 receives as input a "PLL control signal 226" (for ease of discussion, referred to as PLL control signal 226a-n corresponding to clock ratio controller 106a-n).

113. Processor core 110a-n receives respective voltage supply "voltage level signal 138a-n" from voltage controller 107a-n and respective output clock signal "core clock signal 128a-n from clock ratio controller 106a-n for "dynamic voltage and frequency scaled digital systems." *Id.* at [0001]. The full disclosure is provided below:

> The present invention generally relates to digital systems. More particularly, ***it relates to a clock ratio controller for dynamic voltage and frequency scaled digital systems***, and applications thereof.

*Id.* at [0001].

114. Specifically, Knoth discloses that each of voltage controllers 107a-n

receives "individual voltage management (VM) signals 113a-n" that "initiate voltage adjustments" to "individually control/scale the voltage of each processor core 110a-n" so that the voltage supply may increase or decrease. *Id.* at [0025], [0031]-[0032]. Processor core 110a-n thus dynamically receives voltage level 138a-n.

> In an embodiment, power management unit 102 provides individual frequency management (FM) signals 111a-n to clock ratio controllers 106a-n and individual voltage management (VM) signals 113a-n to voltage controllers 107a-n. These signals are also referred to herein as power management signals. A FM signal 111o is also provided to prescaler 104. These FM signals are used, for example, to initiate frequency adjustments. ***The VM signals are used, for example, to initiate voltage adjustments.*** FM signals 111a-n can be used to individually control/scale the frequency of each processor core 110a-n. ***VM signals 113a-n can be used to individually control/scale the voltage of each processor core 110a-n.*** FM signal 111o can be used to control/scale the frequency of all of the components of digital system 100.

*Id.* at [0025].

> ***Voltage controllers 107a-n control the voltage levels used to run processor cores 110a-n.*** As shown in FIG. 1A, in an embodiment, each voltage controller 107a-n sends a voltage ready signal 136a-n to each clock ratio controller 106a-n to acknowledge when a new voltage level has been set. Each voltage controller 107a-n also sends a voltage or voltage level signal 138a-n to each processor core 110a-n to power each processor core 110a-n.
>
> In an embodiment, ***if VM signals 113a-n indicate a voltage decrease that would require a frequency decrease, voltage controllers 107a-n may wait for frequency ready signals 134a-n before adjusting the voltage.*** If FM signals 111a-n indicate a frequency increase that would require a voltage increase, clock ratio controllers 106a-n may wait for voltage ready signals 136a-n before adjusting the frequency.

*Id.* at [0031]-[0032].

115. Knoth also discloses that each of clock ratio controllers 106a-n receives "individual frequency management (FM) signals 111a-n" that "initiate frequency adjustments" to "individually control/scale the frequency of each processor core 110a-n" so that core clock signals 128a-n are dynamically modified. *Id.* at [0025], [0027], [0032]. Processor cores 110a-n thus dynamically receive core clock signals 128a-n.

> In an embodiment, power management unit 102 provides individual frequency management (FM) signals 111a-n to clock ratio controllers 106a-n and individual voltage management (VM) signals 113a-n to voltage controllers 107a-n. These signals are also referred to herein as power management signals. A FM signal 111o is also provided to prescaler 104. ***These FM signals are used, for example, to initiate frequency adjustments.*** The VM signals are used, for example, to initiate voltage adjustments. ***FM signals 111a-n can be used to individually control/scale the frequency of each processor core 110a-n.*** VM signals 113a-n can be used to individually control/scale the voltage of each processor core 110a-n. FM signal 111o can be used to control/scale the frequency of all of the components of digital system 100.

*Id.* at [0025].

> ***Clock ratio controllers 106a-n generate the clock or clocking signals used to control processor cores 110a-n*** and data communications on busses 122a-n. As shown in FIG. 1A, in an embodiment, each clock ratio controller 106a-n generates a core clock signal 128 used to drive an associated processor core 110 and a bus clock signal 120 used to drive an associated transmitter/receiver or transceiver 142 of coherency manager 108. These clocking signals are based on the fast clock signal 114a-n received by clock ratio controllers 106a-n. Core clock signals 128a-n and bus clock signals 120a-n are not required to operate at the

> same frequency, nor do core clock signals 128a-n have to be faster or slower than bus clock signals 120a-n.

*Id.* at [0027].

> In an embodiment, if VM signals 113a-n indicate a voltage decrease that would require a frequency decrease, voltage controllers 107a-n may wait for frequency ready signals 134a-n before adjusting the voltage. ***If FM signals 111a-n indicate a frequency increase that would require a voltage increase, clock ratio controllers 106a-n may wait for voltage ready signals 136a-n before adjusting the frequency.***

*Id.* at [0032].

116. Thus, Knoth discloses that the first set of cores (processor core 110a extendable to include a set of associated cores, or alternatively, processor core 110a of Knoth modified in view of Allarey to include a set of multiple cores) dynamically receives a first supply voltage (voltage level 138a) and a first output clock signal (core clock 128a) of a first phase lock loop (PLL) (PLL 202a) having a first clock signal as input (the timing pulses from oscillator 219a, or alternatively, PLL control signal 226a); the second set of cores (processor core 110n extendable to include a set of associated cores, or alternatively, processor core 110n of Knoth modified in view of Allarey to include a set of multiple cores) dynamically receives a second supply voltage (voltage level 138n) and a second output clock signal (core clock 128n) of a second phase lock loop (PLL) (PLL 202n) having a second clock signal as input (the timing pulses from oscillator 219n, or alternatively, PLL control signal 226n).

117. Allarey also discloses the claim limitation "wherein each processor

core from the first set of processor cores is configured to dynamically receive a first supply voltage and a first output clock signal of a first phase lock loop (PLL) having a first clock signal as input." Specifically, with regards to dynamically receiving a first voltage supply, Allarey discusses that the processor cores in site 0 and the processor cores in site 1 are supplied with voltage from voltage plane 106. Ex[1006] at 2:41-44, 2:59-3:24. "Site 1(104) may have different voltage requirements than site 0 (102) at any given time" and requests a different amount of voltage:

> ***Site 1 (104) may have different voltage requirements than site 0 (102) at any given time.*** Thus, in many embodiments, site 1 (104) communicates its needed voltage to site 0 (102) across the PMLink 120 so ***site 0 (102) can request at least that amount of voltage from the voltage regulator 126***.

*Id.* at 3:20-24.

118. Further, Allarey discloses that the processor cores in site 0 (for example, cores 108 and 110 in site 0 (102)) receive the clock signal generated by the PLL for site 0 (PLL 116 for site 0 (102)) (highlighted in purple). The processor cores in site 1 (for example, cores 112 and 114 in site 1 (104)) receive the clock signal generated by the PLL for site 1 (PLL 118 for site 1 (104)) (highlighted in purple).

> FIG. 1 is an illustration of an apparatus to stabilize a supplied voltage during a clock signal frequency locking process according to some embodiments.
>
> A quad-core processor with two dual-core dies 100 is portrayed. In many embodiments, the processor 100 includes two sites, site 0 (102) and site 1 (104). Both sites are coupled to a common voltage plane 106. ***Site 0 (102) includes two processing cores, core 0 (108) and core 1***

> ***(110).*** Site 1 (104) also includes two processing cores, core 0 (112) and core 1 (114). Each core includes logic to execute instructions. Thus, combined, the two sites have a total of four cores, hence the processor being a quad-core processor with two dual-core dies. ***Each site includes a phase locked loop (PLL) clock signal generation circuit, PLL 116 for site 0 (102)*** and PLL 118 for site 1 (104). Each PLL is capable of generating a clock signal that the cores located at each respective site can use as a reference clock. Additionally, each PLL can change the frequency of the clock signal through a relocking process. In other embodiments that are not shown, an alternative form of clock signal generation logic generates the clock signal supplied to the cores within the processor 100.
>
> In many embodiments, a power management link (PMLink) 120 communicatively couples site 0 and site 1. The specific details of the PMLink 120 and its interface to each site can comprise one of many different link (i.e. interconnect, bus) forms. Generally, the PMLink 120 is capable of transmitting data back and forth between site 0 (102) and site 1 (104). In many embodiments, there is a master site and a slave site with respect to controlling the voltage level supplied to the voltage plane 106. In many embodiments, site 0 (102) is capable of controlling the voltage level supplied to the voltage plane 106. The voltage control process may be referred to as voltage correction.

*Id.* at 2:38-3:3.

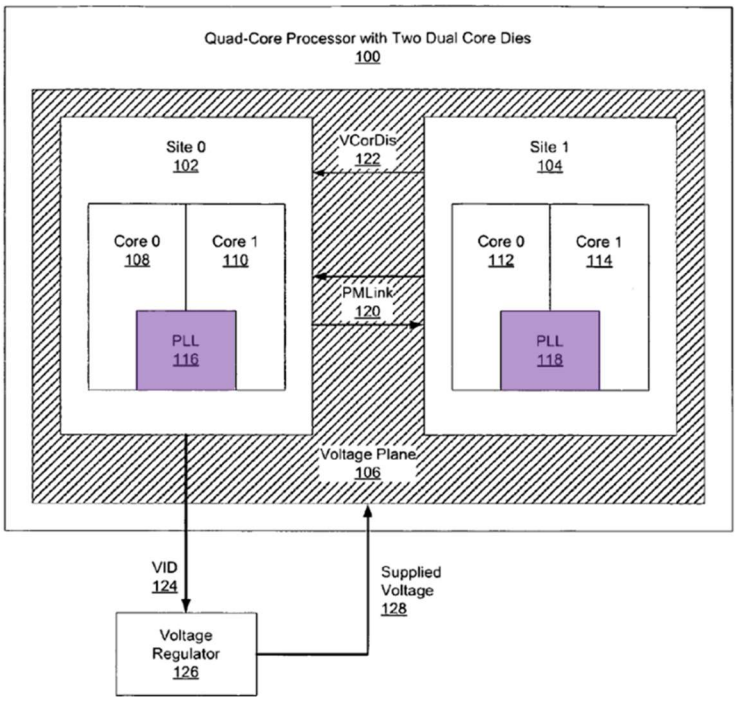


*FIG. 1 of Allarey, annotated*

119. Allarey further states that "[t]he voltage supplied to the processor and the frequency of the processor may be dynamically modified." *Id.* at 1:15-17. Allarey discloses "dynamically modify[ing]" "[t]he voltage supplied to the processor and the frequency of the processor." *Id.* at 1:12-24. Specifically, Allarey discloses "dynamically modify[ing] a voltage supplied to the first site" and "voltage regulator 126 … regulates the supplied voltage 128 … to supply to the voltage plane 106." *Id.* at Abstract and 3:4-19, FIG. 1 (reproduced above, annotated). Accordingly, Allarey discloses that the cores in site 0 and site 1 dynamically receive respective supply voltages that vary as adjusted by circuitry of the voltage regulator 126:

> A modern multi-core processor, such as an Intel® architecture processor or another brand processor, generally has multiple power

> states available to allow for power conservation when the processor is not busy. ***The voltage supplied to the processor and the frequency of the processor may be dynamically modified during operation based on a number of factors such as the current power state of the processor.*** It is generally beneficial to have a stable and unchanging voltage supplied to a clock signal generation circuit, such as a phase locked loop (PLL), when the PLL is in the process of modifying (e.g. relocking) the frequency of the clock signal being output. Asynchronous voltage changes during this time may disrupt a PLL lock process.

*Id.* at 1:12-24.

> A processor, system, and method are disclosed. In an embodiment, the processor includes a first site and a second site. There is a link to transmit a voltage stabilization signal from the second site to the first site. ***In the first site voltage correction logic can dynamically modify a voltage supplied to the first site and second site.*** In the second site there is logic to assert the voltage stabilization signal. After asserting the voltage stabilization signal, the second site is granted at least a window of time in which the supplied voltage to the second site does not change.

*Id.* at Abstract.

> In many embodiments, logic within site 0 (102) sends a voltage identification (VID) value 124 to a voltage regulator 126 external to the processor. ***The voltage regulator 126 interprets the VID value and based on that information, regulates the supplied voltage 128 to the processor 100. Thus, in many embodiments where site 0 (102) and site 1 (104) are supplied with the same voltage through common voltage plane 106,*** logic within site 0 (102) dictates the supplied voltage to both site 0 (102) and site 1 (104). In many other embodiments, logic within site 0 (102) may send information other than a VID 124 to the voltage regulator 126 for supplied voltage level modifications. The information sent to the voltage regulator 126 can be in any form as long as it informs the voltage regulator 126 of the new voltage to supply to the voltage plane 106.

*Id.* at 3:4-19.

120. For the clock signaling, "[e]ach PLL can change the frequency of the clock signal." *Id.* at 2:41-58. Further, Allarey discloses that "[e]ach site includes a phase locked loop (PLL) clock signal generation circuit" and "[e]ach PLL is capable of generating a clock signal." *Id.* at 2:41-58. Additionally, Allarey states that "[t]he frequency of the processor may be dynamically modified." *Id.* at 1:12-24. Accordingly, Allarey discloses that the cores in site 0 and site 1 dynamically receive respective output clock signals of a PLL.

> A quad-core processor with two dual-core dies 100 is portrayed. In many embodiments, the processor 100 includes two sites, site 0 (102) and site 1 (104). Both sites are coupled to a common voltage plane 106. Site 0 (102) includes two processing cores, core 0 (108) and core 1 (110). Site 1 (104) also includes two processing cores, core 0 (112) and core 1 (114). Each core includes logic to execute instructions. Thus, combined, the two sites have a total of four cores, hence the processor being a quad-core processor with two dual-core dies. ***Each site includes a phase locked loop (PLL) clock signal generation circuit, PLL 116 for site 0 (102) and PLL 118 for site 1 (104). Each PLL is capable of generating a clock signal that the cores located at each respective site can use as a reference clock. Additionally, each PLL can change the frequency of the clock signal through a relocking process.*** In other embodiments that are not shown, an alternative form of clock signal generation logic generates the clock signal supplied to the cores within the processor 100.

*Id.* at 2:41-58.

> A modern multi-core processor, such as an Intel® architecture processor or another brand processor, generally has multiple power states available to allow for power conservation when the processor is not busy. ***The voltage supplied to the processor and the frequency of the processor may be dynamically modified during operation based on a number of factors such as the current power state of the processor.*** It is generally beneficial to have a stable and unchanging

> voltage supplied to a clock signal generation circuit, such as a phase locked loop (PLL), when the PLL is in the process of modifying (e.g. relocking) the frequency of the clock signal being output. Asynchronous voltage changes during this time may disrupt a PLL lock process.

*Id.* at 1:12-24.

121. Accordingly, Allarey discloses that the first set of cores (cores in site 0) dynamically receives a first supply voltage and a first output clock signal of a first PLL (PLL for site 0). Allarey also discloses that the second set of cores (cores in site 1) dynamically receives a second supply voltage and a second output clock signal of a second PLL (PLL for site 1).

**4. 1[b1] - a second set of processor cores of the multi-core processor,**

122. As discussed above for Claim 1[a1] in Ground 1, Knoth and Allarey, alone or in combination, disclose this claim limitation. The disclosure regarding "a first set of processor cores of the multi-core processor" also apply to "a second set of processor cores of the multi-core processor."

**5. 1[b2] - wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage and a second output clock signal of a second PLL having a second clock signal as input,**

123. As explained above for Claim 1[a2] in Ground 1, Knoth and Allarey, alone or in combination, disclose this claim limitation. The disclosure regarding the first set of processor cores of the multi-core processor also apply to the second set of processor cores of the multi-core processor.

### 6. 1[b3] - wherein the first supply voltage is independent from the second supply voltage, and

124. Knoth discloses "wherein the first supply voltage is independent from the second supply voltage." In other words, the first supply voltage is separate and unrelated to the second supply voltage. In Knoth, there is no connection between the two supply voltages.

125. Knoth discloses that the first and second sets of cores receive independent supply voltages. Knoth discloses that "VM signals 113a-n can be used to individually control/scale the voltage of each processor core 110a-n" and "[e]ach voltage controller 107a-n also sends a voltage or voltage level signal 138a-n to each processor core 110a-n to power each processor core 110a-n":

> In an embodiment, power management unit 102 provides individual frequency management (FM) signals 111a-n to clock ratio controllers 106a-n and individual voltage management (VM) signals 113a-n to voltage controllers 107a-n. These signals are also referred to herein as power management signals. A FM signal 111o is also provided to prescaler 104. These FM signals are used, for example, to initiate frequency adjustments. ***The VM signals are used, for example, to initiate voltage adjustments.*** FM signals 111a-n can be used to individually control/scale the frequency of each processor core 110a-n. ***VM signals 113a-n can be used to individually control/scale the voltage of each processor core 110a-n.*** FM signal 111o can be used to control/scale the frequency of all of the components of digital system 100.

Ex[1005] at [0025].

> Voltage controllers 107a-n control the voltage levels used to run processor cores 110a-n. As shown in FIG. 1A, in an embodiment, each voltage controller 107a-n sends a voltage ready signal 136a-n to each

> clock ratio controller 106a-n to acknowledge when a new voltage level has been set. ***Each voltage controller 107a-n also sends a voltage or voltage level signal 138a-n to each processor core 110a-n to power each processor core 110a-n.***

*Id.* at [0031].

126. The VM signals are voltage management signals that are sent to the voltage controllers control/scale the voltage. As a result, Knoth teaches individually controlling each voltage level signal 138a-n such that the first supply voltage is independent from the second supply voltage.

127. Allarey also discloses "wherein the first supply voltage is independent from the second supply voltage." Allarey discloses that the first and second sets of cores receive independent supply voltages:

> ***Site 1 (104) may have different voltage requirements than site 0 (102) at any given time.*** Thus, in many embodiments, site 1 (104) communicates its needed voltage to site 0 (102) across the PMLink 120 so site 0 (102) can request at least that amount of voltage from the voltage regulator 126.

Ex[1006] at 3:20-24.

128. Site 1 (104) having different voltage requirements than site 0 (102) means that the supply voltage to each is different and independent. As a result, Allarey teaches different supply voltage requirements for the sets of processor cores such that the first supply voltage is independent from the second supply voltage.

**7. 1[b4] - the first clock signal is independent from the second clock signal; and**

129. Knoth discloses that "the first clock signal is independent from the

second clock signal." In other words, the first clock signal is separate and unrelated to the second clock signal. In Knoth, there is no connection between the two clock signals.

130. Knoth discloses that clock ratio controllers 106a-n receive, respectively, "individual frequency management (FM) signals 111a-n" that are "used to individually control/scale the frequency of each processor core 110a-n." Ex[1005] at [0025]. Each FM signal is determined based on the needs of each processor core. For example, the first clock signal "operates at a first rate," and the second clock signal "operates at a second rate. *Id.* at [0004]. The first rate and the second rate are not equal and do not rely on each other. Thus, clock ratio controllers 106a-n operate independently from each other. As such, the first clock signal of the first PLL within clock ratio controller 106a is independent from the second clock signal of the second PLL within clock ratio controller 106n. *Id*.

> In an embodiment, ***power management unit 102 provides individual frequency management (FM) signals 111a-n to clock ratio controllers 106a-n*** and individual voltage management (VM) signals 113a-n to voltage controllers 107a-n. These signals are also referred to herein as power management signals. A FM signal 111o is also provided to prescaler 104. These FM signals are used, for example, to initiate frequency adjustments. The VM signals are used, for example, to initiate voltage adjustments. ***FM signals 111a-n can be used to individually control/scale the frequency of each processor core 110a-n.*** VM signals 113a-n can be used to individually control/scale the voltage of each processor core 110a-n. FM signal 111o can be used to control/scale the frequency of all of the components of digital system 100.

*Id.* at [0025].

> ***In an embodiment, the present invention provides a digital system that includes a first digital circuit that operates at a first rate determined by a first clock signal and a second digital circuit that operates at a second rate determined by a second clock signal.*** The first digital circuit is coupled to the second digital circuit by a bus that is used for communicating between the first digital circuit and the second digital circuit. A clock ratio controller is used to adjust the frequency of the first clock signal and/or the second clock signal in response to a power management signal without causing a loss of synchronization between the first digital circuit and the second digital circuit.

*Id.* at [0004].

131. Furthermore, as explained above for Claim 1[a2] in Ground 1, Knoth discloses that each clock ratio controller 106 each includes a phase locked loop (PLL) 202 having as input "timing pulses generated by an oscillator 219" within the corresponding clock ratio controller 106:

> FIGS. 2A-B are diagrams that further illustrate an example clock ratio controller 106 according to an embodiment of the present invention. As shown in FIG. 2A, ***clock ratio controller 106 is coupled to a phase-locked loop (PLL) 202*** and prescaler 104 of digital system 100.
>
> ***PLL 202 outputs a clock signal 218 based on, for example, timing pulses generated by an oscillator 219.*** Clock signal 218 is provided to prescaler 104 and a frequency controller 204. In an embodiment, PLL 202 is controlled by frequency controller 204 using a PLL control signal 226.

*Id.* at [0041]-[0042].

132. A PHOSITA would read Knoth to include independent first and second PLLs within clock ratio controller 106a and 106n, respectively, receiving the first and second clock signals (the timing pulses generated by oscillators 219a-n within

clock ratio controller 106a and 106n, respectively) that are independent from each other. Each of the PLLs 202 receive clock signals from oscillators 219 that are independent of each other.

133. Alternatively, PLL control signals 226a and 226n of Knoth disclose that the first and second clock signals are independent of each other, as claimed. Each PLL control signal 226 is generated based on its respective, independent "voltage ready signal 136" (136a-n) and "FM signal 111" (111a-n). Ex[1005] at [0055]. Specifically, each voltage ready signal 136a-n is provided by a separate voltage controller 107a-n operating independently from each other. *Id.* at [0025], [0031]. Knoth also discloses that "power management unit 102 provides individual frequency management (FM) signals 111a-n" that separate from each other. *Id.* at [0025]. As a result, the PLL control signals 226a and 226n are different and independent from each other.

> In an embodiment, programmable logic array 250 implements boolean logic to perform the below mentioned functions. Programmable logic array 250 generates resync request signal 130, frequency ready signal 134, core fast signal 212, strobe (fast-to-slow) signal 214, strobe (slow-to-fast) signal 216, ratio signals 224a and 224b, and ***PLL control signal 226 based on voltage ready signal 136,*** fast clock frequency signal 220, slow clock frequency signal 222, ***an FM signal 111***, and a count value 264.

*Id.* at [0055].

> In an embodiment, ***power management unit 102 provides individual frequency management (FM) signals 111a-n to clock ratio controllers 106a-n and individual voltage management (VM) signals 113a-n to***

> ***voltage controllers 107a-n.*** These signals are also referred to herein as power management signals. A FM signal 111o is also provided to prescaler 104. These FM signals are used, for example, to initiate frequency adjustments. The VM signals are used, for example, to initiate voltage adjustments. ***FM signals 111a-n can be used to individually control/scale the frequency of each processor core 110a-n. VM signals 113a-n can be used to individually control/scale the voltage of each processor core 110a-n.*** FM signal 111o can be used to control/scale the frequency of all of the components of digital system 100.

*Id.* at [0025].

> ***Voltage controllers 107a-n control the voltage levels used to run processor cores 110a-n.*** As shown in FIG. 1A, in an embodiment, ***each voltage controller 107a-n sends a voltage ready signal 136a-n to each clock ratio controller 106a-n to acknowledge when a new voltage level has been set.*** Each voltage controller 107a-n also sends a voltage or voltage level signal 138a-n to each processor core 110a-n to power each processor core 110a-n.

*Id.* at [0031].

134. Therefore, Knoth discloses that the first clock signal is independent from the second clock signal.

135. Allarey also discloses "the first clock signal is independent from the second clock signal." In Allarey, each site includes a PLL clock signal generation circuit:

> A quad-core processor with two dual-core dies 100 is portrayed. In many embodiments, the processor 100 includes two sites, site 0 (102) and site 1 (104). Both sites are coupled to a common voltage plane 106. Site 0 (102) includes two processing cores, core 0 (108) and core 1 (110). Site 1 (104) also includes two processing cores, core 0 (112) and core 1 (114). Each core includes logic to execute instructions. Thus, combined, the two sites have a total of four cores, hence the processor being a quad-core processor with two dual-core dies. ***Each site includes***

> ***a phase locked loop (PLL) clock signal generation circuit, PLL 116 for site 0 (102) and PLL 118 for site 1 (104). Each PLL is capable of generating a clock signal that the cores located at each respective site can use as a reference clock. Additionally, each PLL can change the frequency of the clock signal through a relocking process.*** In other embodiments that are not shown, an alternative form of clock signal generation logic generates the clock signal supplied to the cores within the processor 100.

Ex[1006] at 2:41-58.

136. The PLL clock signal generation circuits are independent of each other since "each PLL can change the frequency of the clock signal through a relocking process." *Id*. When the PLL frequency is changed, the PLL clock signal generation circuit is also modified. Therefore, the first PLL clock signal generation circuit is independent from the second PLL clock signal generation circuit.

**8. 1[c1] - an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores,**

137. Knoth teaches "an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores." An interface block coupled to the two sets of processor cores requires an element that is connected to the first set and connected to the second set of processor cores. As depicted in FIG. 1A, Knoth discloses a coherency manager 108 (highlighted in green) coupled to the processor cores 110a-n (processor cores 110a and 110n highlighted in yellow):

> Coherency manager 108 coordinates cache memory access for processor cores 110a-n. As shown in FIG. 1A, ***coherency manager 108 and processor cores 110a-n are coupled together by busses 122a-n.*** In

an embodiment, coherency manager 108 runs at the same clock speed as the fastest processor core 110a-n.

Ex[1005] at [0033].

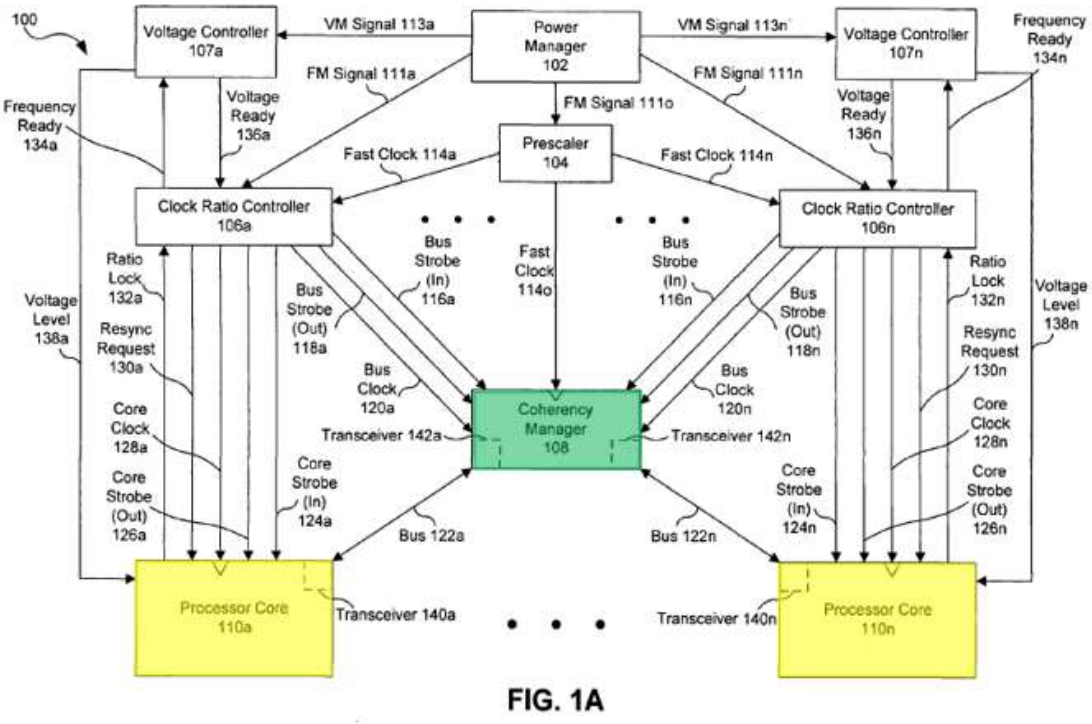


FIG. 1A

*FIG. 1A of Knoth, annotated*

138. The coherency manager 108 and the processor cores 110a-n are coupled together by busses 122a-n. *Id.* at [0033]. Knoth at FIG. 1B shows the details of the coupling between the coherency manager 108 and the processor core 110.

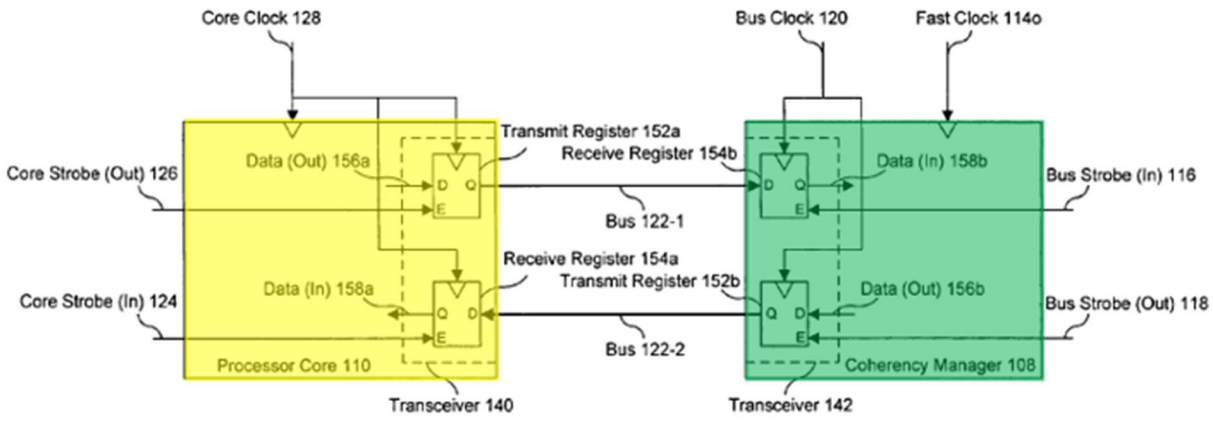


*FIG. 1B of Knoth, annotated*

139. FIG. 1B of Knoth illustrates that the coherency manager 108 is connected to a representative processor core 110 via busses 122-1 and 122-2. Bus 122-1 transmits data from the processor core 110 to the coherency manager 108, and bus 122-2 transmits data in the other direction.

140. Therefore, Knoth discloses an interface block (coherency manager 108) coupled to the first set of processor cores (110a) and also coupled to the second set of processor cores (110n).

141. Allarey also discloses "an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores." As depicted in FIG. 1, Allarey discloses a power management link (PMLink) 120 (highlighted in green) that couples the first and second sets of cores in site 0 and site 1 (highlighted in yellow). Ex[1006] at 2:59-3:3, FIG. 1A. PMLink 120 is described as an interface to couple site 0 and site 1.

> In many embodiments, ***a power management link (PMLink) 120 communicatively couples site 0 and site 1. The specific details of the PMLink 120 and its interface to each site can comprise one of many different link (i.e. interconnect, bus) forms.*** Generally, the PMLink 120 is capable of transmitting data back and forth between site 0 (102) and site 1 (104). In many embodiments, there is a master site and a slave site with respect to controlling the voltage level supplied to the voltage plane 106. In many embodiments, site 0 (102) is capable of controlling the voltage level supplied to the voltage plane 106. The voltage control process may be referred to as voltage correction.

*Id.* at 2:59-3:3.

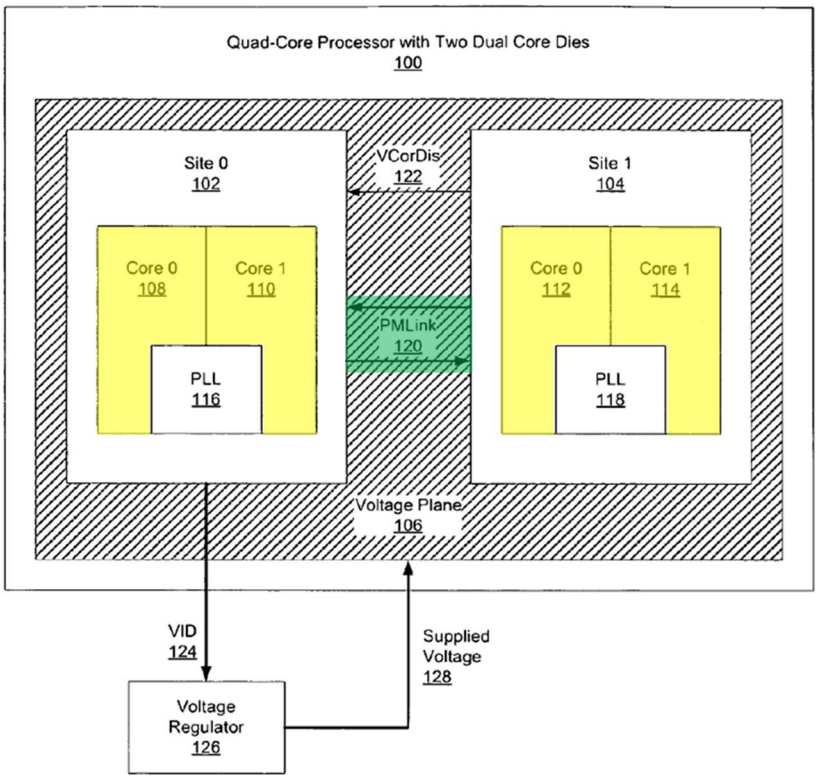


*FIG. 1 of Allarey, annotated*

142. Additionally, Allarey at FIG. 3 illustrates another embodiment that also teaches this claim limitation. FIG. 3 discloses a power management link (PMLink) 328 (highlighted in green) that couples the first and second sets of cores in site 0 and

site 1 (highlighted in yellow). Ex[1006] at 5:34-43, FIG. 3. PMLink 328 is described as an interface to couple site 0 and site 1. The details of Allarey are provided below.

> In many embodiments, ***a power management link (PMLink) 328 communicatively couples site 0 and site 1. The specific details of the PMLink 328 and its interface to each site are discussed above in reference to FIG. 1.*** In many embodiments, there is a master site and a slave site with respect to controlling the voltage level supplied to the voltage plane 306. In many embodiments, logic within site 0 (302) is capable of controlling the voltage level supplied to the voltage plane 306. The voltage control process may be referred to as voltage correction.

*Id.* at 5:34-43.

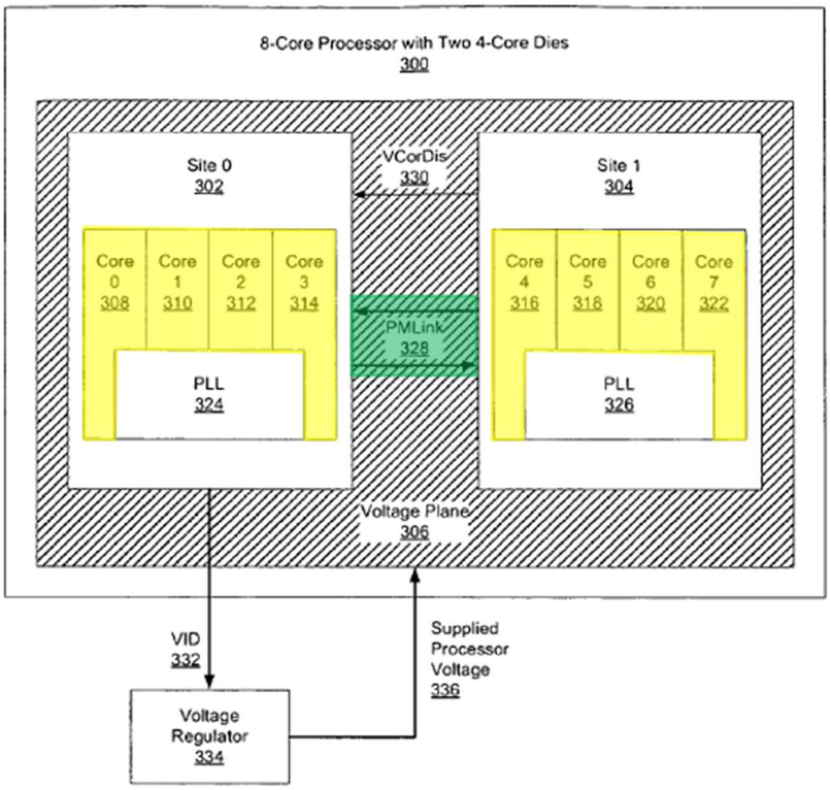


*FIG. 3 of Allarey, annotated*

143. Therefore, Allarey discloses an interface block (PMLink 120 or 328) coupled to the first set of processor cores (cores 108 and 110 or cores 308-314) and

also coupled to the second set of processor cores (cores 112 and 114 or cores 316-322).

**9. 1[c2] - wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores.**

144. Knoth discloses "wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores." Knoth discloses that the coherency manager 108 "facilitate[s] communications between the first digital circuit and the second digital circuit." Ex[1005] at [0006]. Further, the coherency manager 108 has "a bus that is used for communications between the first digital circuit and the second digital circuit." *Id.* at Abstract. As a result, Knoth teaches that the interface block (coherency manager 108) is configured to facilitate communication between the first set of processor cores 110a and the second set of processor cores 110n.

> In an embodiment, the clock ratio controller outputs strobe signals that enable operation of transmit registers and receive registers of the first digital circuit and the second digital circuit. ***These transmit registers and receive registers are used to facilitate communications between the first digital circuit and the second digital circuit.***

*Id.* at [0006].

> The present invention provides a clock ratio controller for dynamic voltage and frequency scaled digital systems, and applications thereof. In an embodiment, a digital system is provided that includes a first digital circuit that operates at a first rate determined by a first clock signal and a second digital circuit that operates at a second rate determined by a second clock signal. ***The first digital circuit is coupled***

> ***to the second digital circuit by a bus that is used for communications between the first digital circuit and the second digital circuit.*** A clock ratio controller is used to adjust the frequency of the first clock signal and/or the second clock signal in response to a power management signal without causing a loss of synchronization between the first digital circuit and the second digital circuit.

*Id.* at Abstract.

145. Allarey also discloses "wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores." Allarey discloses that the "power management link (PMLink)" "communicatively couples site 0 and site 1" and "transmit[s] data back and forth between site 0 (102) and site 1 (104)":

> In many embodiments, ***a power management link (PMLink) 120 communicatively couples site 0 and site 1.*** The specific details of the PMLink 120 and its interface to each site can comprise one of many different link (i.e. interconnect, bus) forms. Generally, ***the PMLink 120 is capable of transmitting data back and forth between site 0 (102) and site 1 (104).*** In many embodiments, there is a master site and a slave site with respect to controlling the voltage level supplied to the voltage plane 106. In many embodiments, site 0 (102) is capable of controlling the voltage level supplied to the voltage plane 106. The voltage control process may be referred to as voltage correction.

Ex[1006] at 2:59-3:3.

> In many embodiments, ***a power management link (PMLink) 328 communicatively couples site 0 and site 1.*** The specific details of the PMLink 328 and its interface to each site are discussed above in reference to FIG. 1. In many embodiments, there is a master site and a slave site with respect to controlling the voltage level supplied to the voltage plane 306. In many embodiments, logic within site 0 (302) is capable of controlling the voltage level supplied to the voltage plane 306. The voltage control process may be referred to as voltage correction.

*Id.* at 5:34-43.

146. Consequently, Allarey discloses that the interface block (PMLink 120 or 328) is configured to facilitate communication between the first set of processor cores (cores 108 and 110 or 308-314) and the second set of processor cores (cores 112 and 114 or 316-322).

147. Accordingly, Knoth and Allarey, alone or in combination, disclose this final limitation, thus rendering Claim 1 obvious.

### C. Dependent Claim 5

#### 1. 5[pre] - The multi-core processor of claim 1,

148. As I explained above for Claim 1, Knoth and Allarey, alone or in combination, teach each and every limitation of Claim 1, rendering it obvious.

#### 2. 5[a] - wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor.

149. Knoth discloses "wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor." The '339 Patent provides no definition or explanation of the claimed "control signals."

150. Knoth discloses "power management unit 102 provides individual frequency management (FM) signals 111a-n" and "individual voltage management (VM) signals 113a-n." Ex[1005] at [0025]. The FM signals 111a-n "are used … to

initiate frequency adjustments" and "to individually control/scale the frequency of each processor core 110a-n." *Id.* The VM signals 113a-n "are used … to initiate voltage adjustments" and "to individually control/scale the Voltage of each processor core 110a-n." *Id.* As such, the power management unit 102 teach the claimed "one or more control blocks."

> In an embodiment, ***power management unit 102 provides individual frequency management (FM) signals 111a-n to clock ratio controllers 106a-n and individual voltage management (VM) signals 113a-n to voltage controllers 107a-n.*** These signals are also referred to herein as power management signals. A FM signal 111o is also provided to prescaler 104. ***These FM signals are used, for example, to initiate frequency adjustments. The VM signals are used, for example, to initiate voltage adjustments. FM signals 111a-n can be used to individually control/scale the frequency of each processor core 110a-n. VM signals 113a-n can be used to individually control/scale the voltage of each processor core 110a-n.*** FM signal 111o can be used to control/scale the frequency of all of the components of digital system 100.

*Id.* at [0025].

151. As depicted below, power management unit 102 (highlighted in red) is located in the periphery of Knoth's multi-core processor. *Id.* at FIG. 1A (annotated). The center of the microprocessor digital system 100 in FIG. 1A of Knoth can be equated to the coherency manager 108. The power management unit 102 is located at a position above and away from the coherency manager 108.

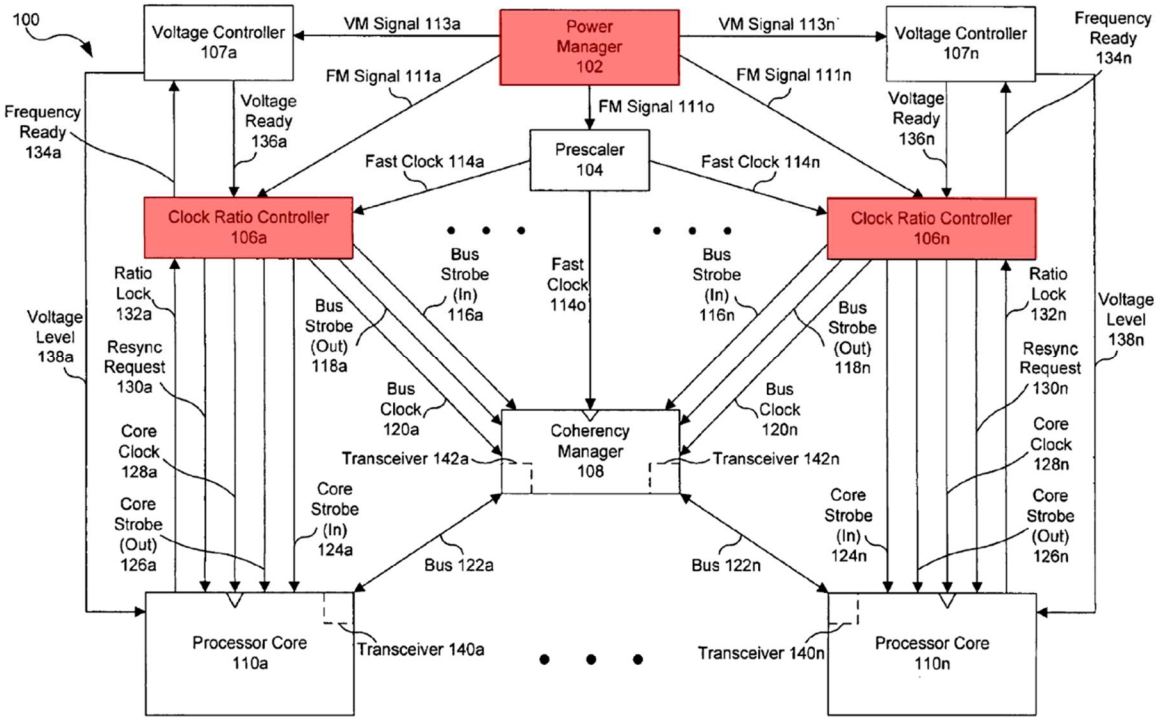


*FIG. 1A of Knoth, annotated*

152. In addition, Knoth discloses that "clock ratio controllers 106a-n … generate resynchronization (resync) requests 130a-n" that coordinates synchronization between the processor cores and controls timing of synchronization and frequency adjustments. *Id.* at [0029]. "[R]esynchronization requests 130a-n are an indication from clock ratio controllers 106a-n to processor cores 110a-n that a clock ratio adjustment is about to occur." *Id.* As depicted above clock ratio controllers 106a-n (highlighted in red) are located in the periphery of Knoth's multi-core processor. The center of the microprocessor digital system 100 in FIG. 1A of Knoth can be equated to the coherency manager 108. The clock ratio controllers 106a-n are located at a position away from the coherency manager 108. Accordingly,

the clock ratio controllers teach the claimed one or more control blocks.

> In an embodiment, ***clock ratio controllers 106a-n also generate resynchronization (resync) requests 130a-n*** and receive ratio lock signals 132a-n. These signals are used, for example, to perform handshaking between clock ratio controllers 106a-n and processor cores 110a-n. In an embodiment, ***resynchronization requests 130a-n are an indication from clock ratio controllers 106a-n to processor cores 110a-n that a clock ratio adjustment is about to occur.*** Ratio lock signals 132a-n are an indication from processor cores 110a-n to clock ratio controllers 106a-n that processor cores 110a-n are ready for a new clock frequency signal.

*Id.* at [0029].

153. The ’339 Patent provides no discussion or indication that the location of the control blocks, whether periphery, centrally or otherwise located relative to the multi-core processor, would in any way limit or alter their function in providing control signals. At most, the ’339 Patent discloses that “the power control block 108 and the clock control block 110 may be arranged at two different sides of the multi-core processor 100” or “at the same side,” which is contrasted with other implementations where “the power control block 108 and the clock control block 110 may be arranged in a common area located near the center of the multi-core processor 100”:

> The multi-core processor 100 may be further divided into regions. In some implementations, the regions of multi-core processor 100 may correspond to rows of the two-dimensional array, and the regions may or may not be overlapping. Each row of processors may also be referred to as a “stripe.” For example, the multi-core processor 100 may be divided into stripes 112, 114, 116, and 118. Each stripe may be associated with an independent power profile. For example, the stripe

> 112 may be powered by a supply voltage received from a power control block 108 and/or may be associated with an independent clock domain defined by a clock signal received from a clock control block 110. In some implementations, the power control block 108 and the clock control block 110 may be arranged at two different sides of the multi-core processor 100 as shown in FIG. 1. ***In some other implementations, the power control block 108 and the clock control block 110 may be arranged at the same side of the multi-core processor 100. In yet some other implementations, the power control block 108 and the clock control block 110 may be arranged in a common area located near the center of the multi-core processor 100.***

Ex[1001] at 2:20-40.

154. Though this recitation teaches the location of the one or more control blocks, there is no explanation as to how the positioning would affect their functionality. A PHOSITA would have understood that to place a control block in the periphery, centrally or in other locations serves the same function of providing control signals for managing or coordinating the operations of processor cores. Positioning a control block at a particular location in a multi-core processor is a matter of design choice, or more likely, space constraints. A PHOSITA would have also understood that this function depends on the logical connections and signal routing, rather than the physical location of the control block relative to the processor. When there are connections to the processor such that control signals can be transmitted, the control block can be positioned at any location relative to the processor. Thus, a control block could be relocated from a substantially central location to the periphery (or vice versa) without affecting its ability to perform this

function. A PHOSITA would have understood that the physical location of a control block would have been driven by routine design preferences such as minimizing of the length of critical paths considering the relative locations of other components of the processor. Therefore, a PHOSITA reading Knoth would have found it obvious that the control blocks of Knoth could be placed at a periphery of the multi-core processor. Therefore, Knoth teaches this claim limitation.

155. Allarey also discloses "wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor."

156. Allarey discloses that "logic within site 0 (102) dictates the supplied voltage to both site 0 (102) and site 1 (104)" by "send[ing] a voltage identification (VID) value 124 to a voltage regulator 126 external to the processor." Ex[1006] at 3:4-19. "The voltage regulator 126 interprets the VID value and based on that information, regulates the supplied voltage 128 to the processor 100." *Id.* Allarey discloses that voltage identification (VID) value 124 "informs … the new voltage to supply" "for supplied voltage modifications." *Id.*

> In many embodiments, a power management link (PMLink) 120 communicatively couples site 0 and site 1. The specific details of the PMLink 120 and its interface to each site can comprise one of many different link (i.e. interconnect, bus) forms. Generally, the PMLink 120 is capable of transmitting data back and forth between site 0 (102) and site 1 (104). In many embodiments, ***there is a master site and a slave site with respect to controlling the voltage level supplied to the voltage plane 106.*** In many embodiments, site 0 (102) is capable of controlling

> the voltage level supplied to the voltage plane 106. The voltage control process may be referred to as voltage correction.

*Id.* at 2:59-3:3.

> In many embodiments, ***logic within site 0 (102) sends a voltage identification (VID) value 124 to a voltage regulator 126 external to the processor.*** The voltage regulator 126 interprets the VID value and based on that information, regulates the supplied voltage 128 to the processor 100. Thus, in many embodiments where site 0 (102) and site 1 (104) are supplied with the same voltage through common voltage plane 106, logic within site 0 (102) dictates the supplied voltage to both site 0 (102) and site 1 (104). In many other embodiments, ***logic within site 0 (102) may send information other than a VID 124 to the voltage regulator 126 for supplied voltage level modifications. The information sent to the voltage regulator 126 can be in any form as long as it informs the voltage regulator 126 of the new voltage to supply to the voltage plane 106.***

*Id.* at 3:4-19.

157. A PHOSITA reading Allarey would have understood that voltage regulator 126 is responsible for adjusting the voltage supplied to the processor cores and to carry out that function, voltage regulator 126 must internally generate one or more control signals that adjust the output voltage level. Such internal control signals are inherently part of how a voltage regulator operates. A voltage regulator controls the voltage level by sending control signals to modify the supply voltage. After receiving the VID value, the voltage regulator will internally send signals to adjust the supplied voltage.

158. Therefore, Allarey teaches that the processor cores of Allarey receive one or more control signals including VID value 124 from logic within site 0 (102)

(highlighted in red) and the control signals from voltage regulator 126 (highlighted in red). Ex[1006], FIG. 1 (annotated). As depicted below, voltage regulator 126 is located at the periphery of the multi-core processor of Allarey. *Id.*; *see also id.* at 3:5-6 (voltage regulator 126 is located "external to" the processor dies).

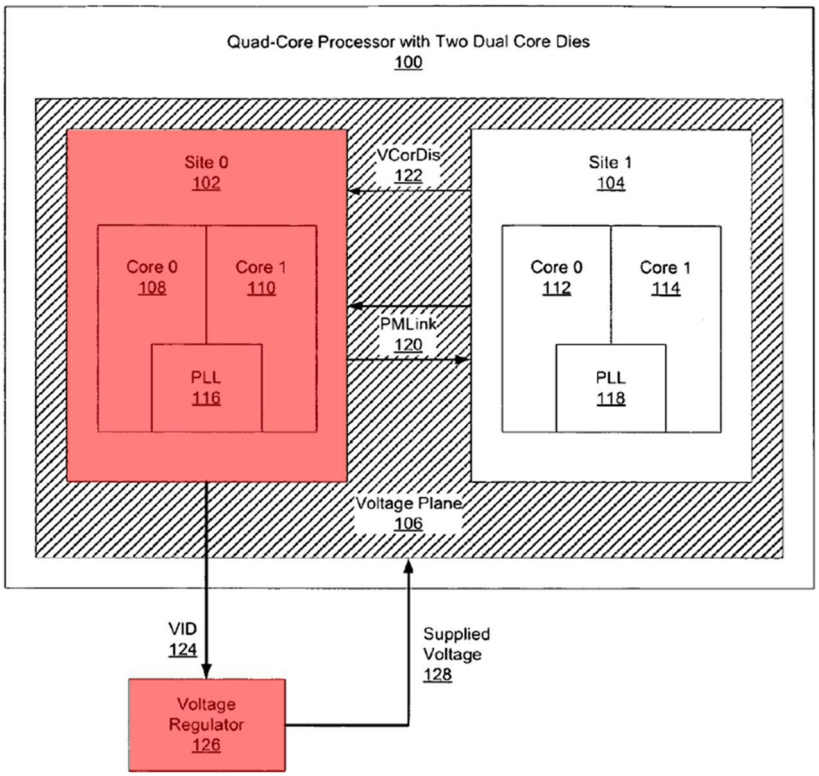


*FIG. 1 of Allarey, annotated*

159. Furthermore, as explained above, a PHOSITA reading Allarey would have found it obvious that the control blocks including voltage regulator 126 could have been placed at a periphery of the multi-core processor without affecting its function of providing control signals to the processor cores. Therefore, Allarey teaches this claim limitation.

160. Therefore, Knoth and Allarey, alone or in combination, teach each and

every limitation of Claim 5, rendering it obvious.

### D. Dependent Claim 8

#### 1. 8[pre] - The multi-core processor of claim 1,

161. As explained above for Claim 1, Knoth and Allarey, alone or in combination, teach each and every limitation of Claim 1, rendering it obvious.

#### 2. 8[a] - wherein the first set of processor cores are located in a first region of the multi-core processor, and the second set of processor cores are located in a second region of the multi-core processor.

162. The '339 Patent discloses that a "region" is simply a spatial grouping of cores based on their physical location in the multi-core processor system:

> ***The multi-core processor 100 may be further divided into regions. In some implementations, the regions of multi-core processor 100 may correspond to rows of the two-dimensional array, and the regions may or may not be overlapping.*** Each row of processors may also be referred to as a "stripe." For example, the multi-core processor 100 may be divided into stripes 112, 114, 116, and 118. Each stripe may be associated with an independent power profile. For example, the stripe 112 may be powered by a supply voltage received from a power control block 108 and/or may be associated with an independent clock domain defined by a clock signal received from a clock control block 110. In some implementations, the power control block 108 and the clock control block 110 may be arranged at two different sides of the multi-core processor 100 as shown in FIG. 1. In some other implementations, the power control block 108 and the clock control block 110 may be arranged at the same side of the multi-core processor 100. In yet some other implementations, the power control block 108 and the clock control block 110 may be arranged in a common area located near the center of the multi-core processor 100.

Ex[1001] at 2:20-40.

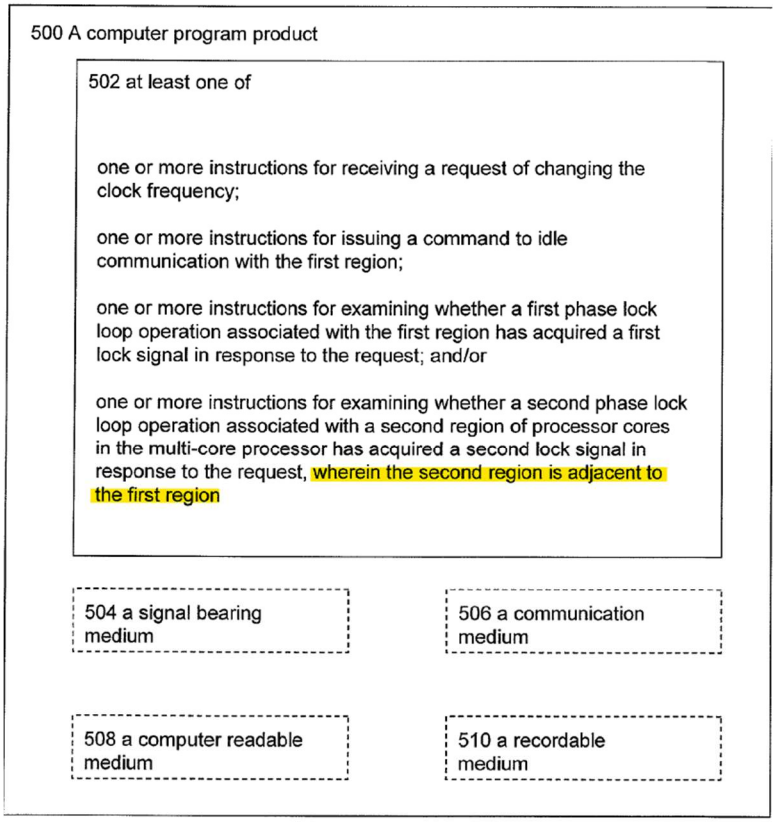


FIG. 5

Ex[1001] at Fig. 5, annotated.

163. Knoth discloses that the multi-core processor system comprises an arbitrary number of cores a through n, including at least two sets of processor cores, for example, 110a and the associated cores (e.g., cores b-i) (highlighted in yellow) and 110n and the associated cores (e.g., cores j-n) (highlighted in orange). Knoth discloses "the first set of cores are located in a first region" and "the second set of processor cores are located in a second region" of the multi-core processor.

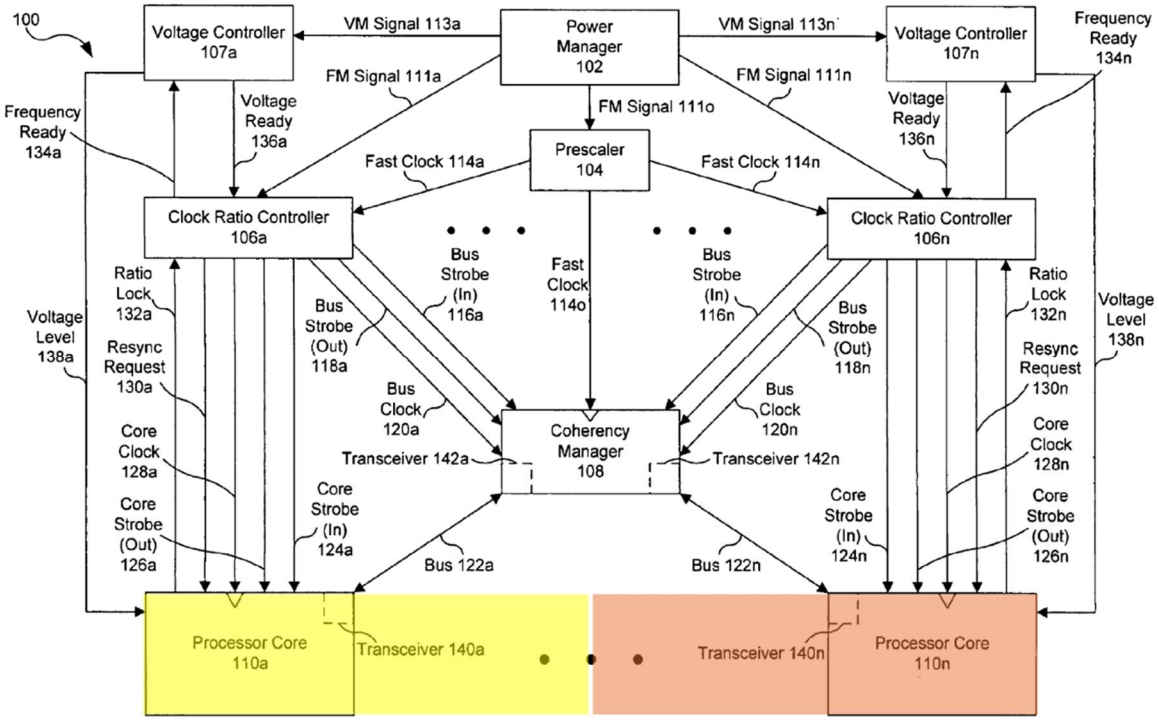


*FIG. 1A of Knoth, annotated*

164. The first set of cores (core 110a and associated cores 110b-i) are located, for instance, in the physical space depicted below in red dotted box; the second set of cores (core 110n and associated cores 110j-n) of Knoth are located in the physical space depicted below in green dotted box. As such, Knoth includes a first region corresponding to the physical location reflected by the red dotted box, and a second region corresponding to the physical location reflected by the green dotted box.

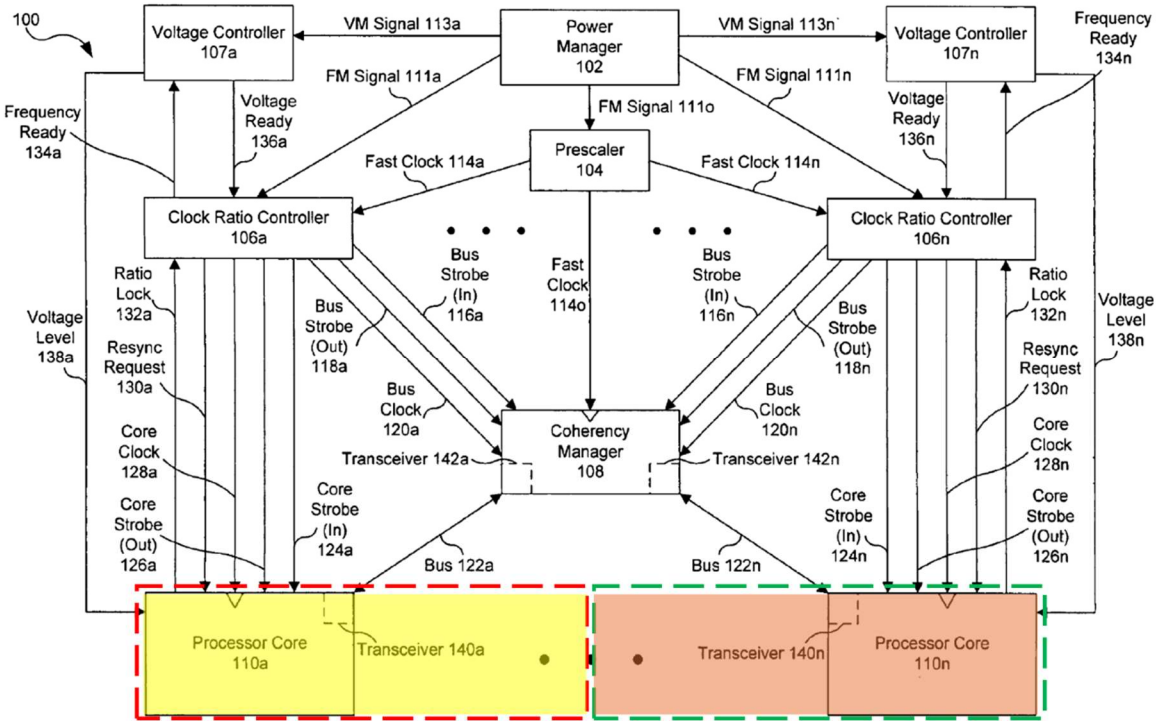


*FIG. 1A of Knoth, annotated*

165. Alternatively, for instance, the first set of cores (core 110a and associated cores 110b-i) are located in the physical space depicted below in purple dotted box; the second set of cores (core 110n and associated cores 110j-n) of Knoth are located in the physical space depicted below in blue dotted box. As such, Knoth includes a first region corresponding to the physical location reflected by the purple dotted box, and a second region corresponding to the physical location reflected by the blue dotted box.

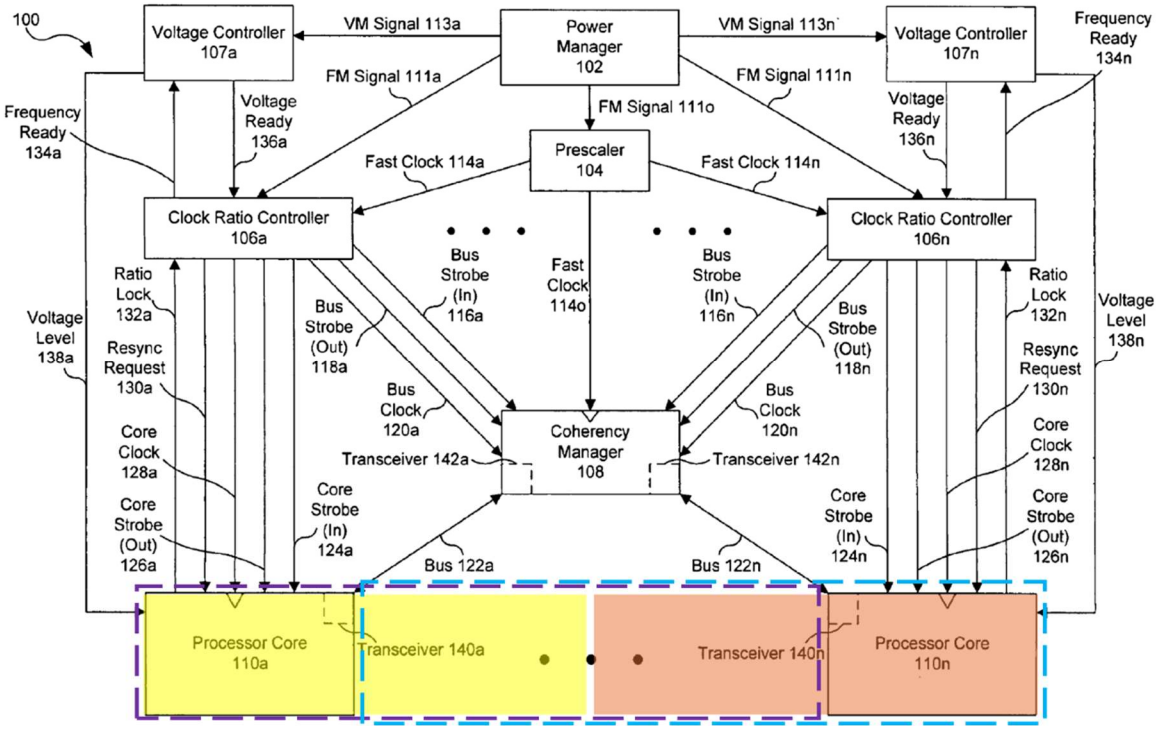


166. Thus, Knoth discloses "wherein the first set of processor cores are located in a first region of the multi-core processor, and the second set of processor cores are located in a second region of the multi-core processor."

167. Allarey discloses that the multi-core processor system comprises multiple cores physically located on two sites. *See, e.g.*, Ex[1006] at 1:52-53, 5:13-33. As discussed above, the multi-core processor of FIG. 3 of Allarey includes the first set of processor cores (for example, core 0 (308), core 1 (310), core 2 (312), and core 3 (314)) and the second set of processor cores (for example, core 4 (316), core 5 (318), core 6 (320), and core 7 (322)). Allarey discloses "the first set of cores are located in a first region" and "the second set of processor cores are located in a second region" of the multi-core processor.

168. For instance, the first set of cores (core 0 (308), core 1 (310), core 2 (312), and core 3 (314)) of Allarey are located in the physical space depicted below in red dotted box; the second set of cores (core 4 (316), core 5 (318), core 6 (320), and core 7 (322)) of Allarey are located in the physical space depicted below in green dotted box. As such, Allarey includes a first region corresponding to the physical location where cores 0-3 are located (in red dotted box), and a second region corresponding to the physical location where cores 4-7 are located (in green dotted box).

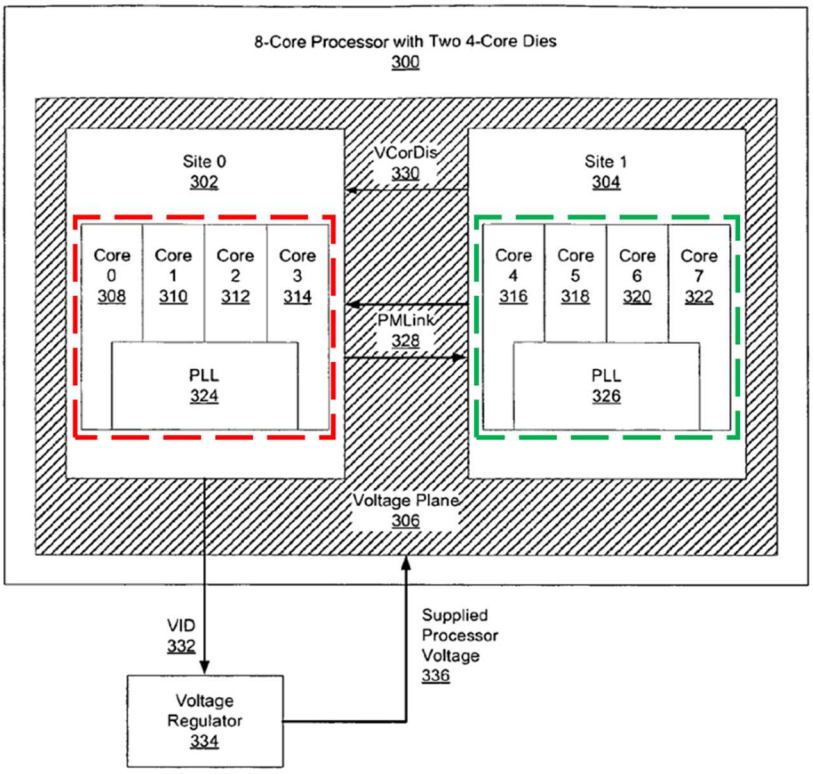


*FIG. 3 of Allarey, annotated*

169. Alternatively, for instance, the first set of cores (core 0 (308), core 1

(310), core 2 (312), core 3 (314)) are located in the physical space depicted below in the purple dotted box; the second set of cores (core 4 (316), core 5 (318), core 6 (320), and core 7 (322)) of Allarey are located in the physical space depicted below in the blue dotted box. As such, Allarey includes a first region where cores 0-6 are located, and a second region where cores 1-7 are located.

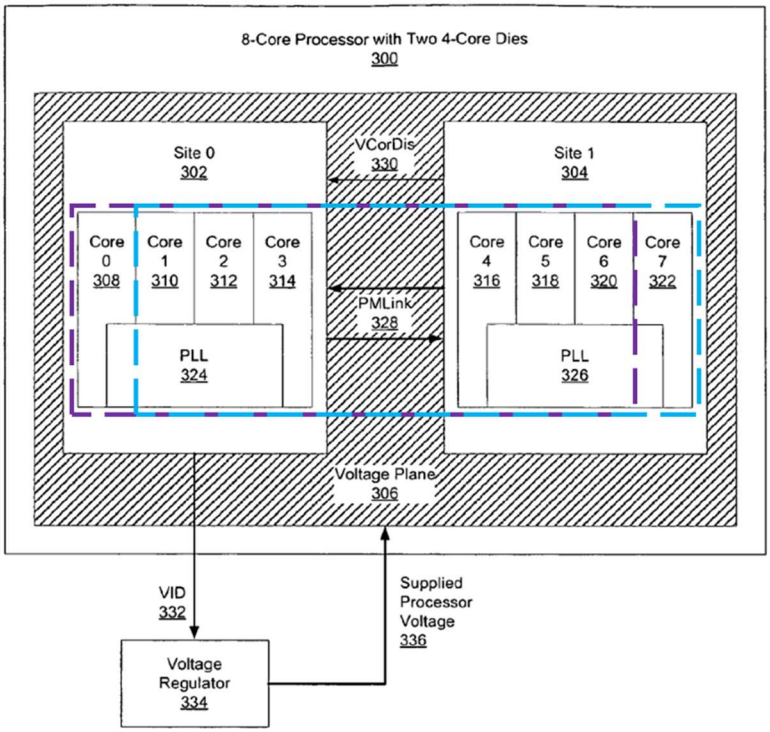


*FIG. 3 of Allarey, annotated*

170. Thus, Allarey discloses "wherein the first set of processor cores are located in a first region of the multi-core processor, and the second set of processor cores are located in a second region of the multi-core processor."

171. Accordingly, Knoth and Allarey, alone or in combination, teach each and every limitation of Claim 8.

### E. Dependent Claim 9

#### 1. 9[pre] - The multi-core processor or claim 8,

172. As explained above for Claim 8, Knoth and Allarey, alone or in combination, teach each and every limitation of Claim 8, rendering it obvious.

#### 2. 9[a] - wherein the first region and the second region are overlapping regions of the multi-core processor.

173. As explained above for Claim 8[a] in Ground 1, Knoth and Allarey, alone or in combination, disclose that the first set of processor cores are located in a first region, and the second set of processor cores are located in a second region, wherein the first region and the second region are overlapping.

174. Again, in Knoth, the first set of cores (core 110a and associated cores 110b-i) are located in the physical space depicted below in purple dotted box; the second set of cores (core 110n and associated cores 110j-n) of Knoth are located in the physical space depicted below in blue dotted box. The first and second regions depicted below share some physical space and thus are overlapping. Accordingly, Knoth discloses this claim limitations.

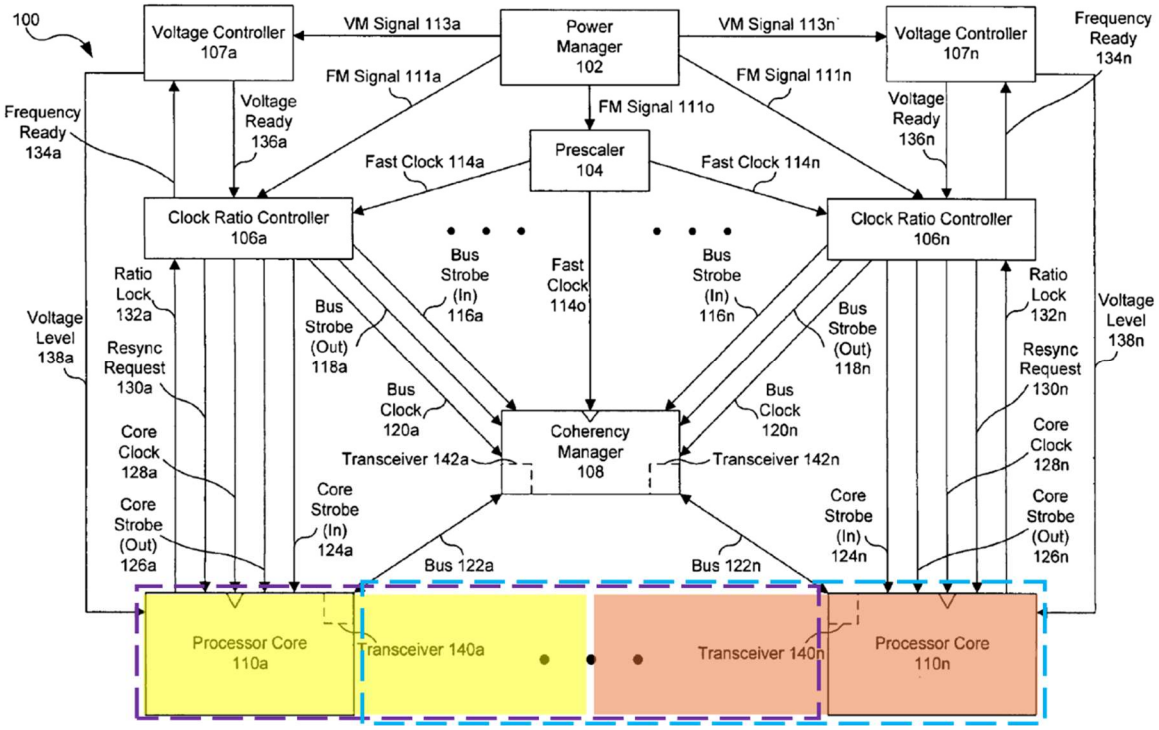


175. In Allarey, the first set of cores (for example, core 0 (308), core 1 (310), core 2 (312), core 3 (314)) are located in the first region depicted below in the purple dotted box; the second set of cores (core 4 (316), core 5 (318), core 6 (320), and core 7 (322)) of Allarey are located in the second region depicted below in the blue dotted box. The first and second regions depicted below share some physical space and thus are overlapping. Accordingly, Allarey discloses this claim limitations.

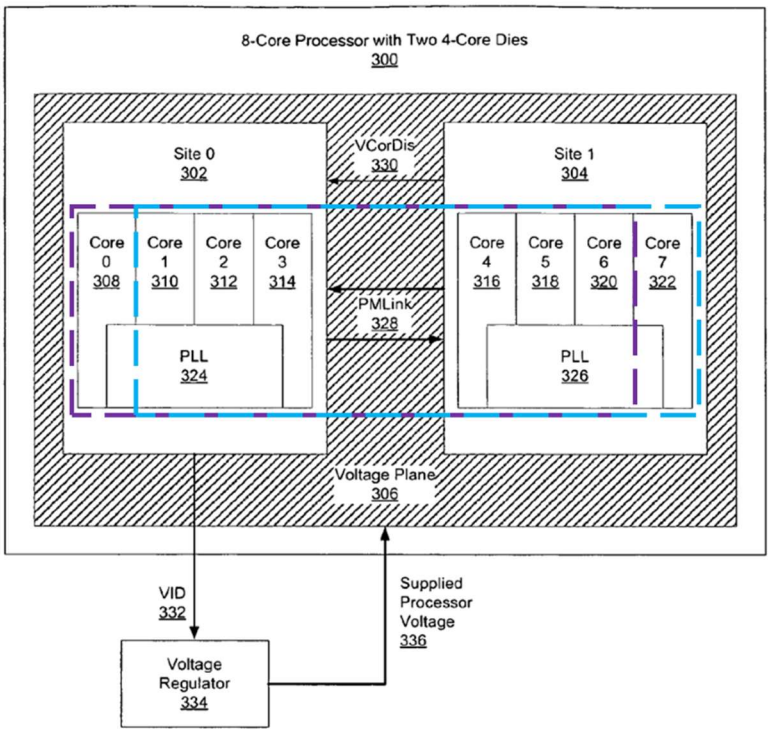


*FIG. 3 of Allarey, annotated*

## F. Dependent Claim 10

### 1. 10[pre] - The multi-core processor of claim 8,

176. As explained above for Claim 8, Knoth and Allarey, alone or in combination, teach each and every limitation of Claim 8, rendering it obvious.

### 2. 10[a] - wherein the first region and the second region are non-overlapping regions of the multi-core processor.

177. As explained above for Claim 8[a] in Ground 1, Knoth and Allarey, alone or in combination, disclose that the first set of processor cores are located in a first region, and the second set of processor cores are located in a second region, wherein the first region and the second region are non-overlapping regions.

178. Again, in Knoth, the first set of cores (core 110a and associated cores

110b-i) are located in the physical space depicted below in red dotted box; the second set of cores (core 110n and associated cores 110j-n) of Knoth are located in the physical space depicted below in green dotted box. The first and second regions depicted below are completely separate in physical layout and thus are non-overlapping. Accordingly, Knoth discloses this claim limitation.

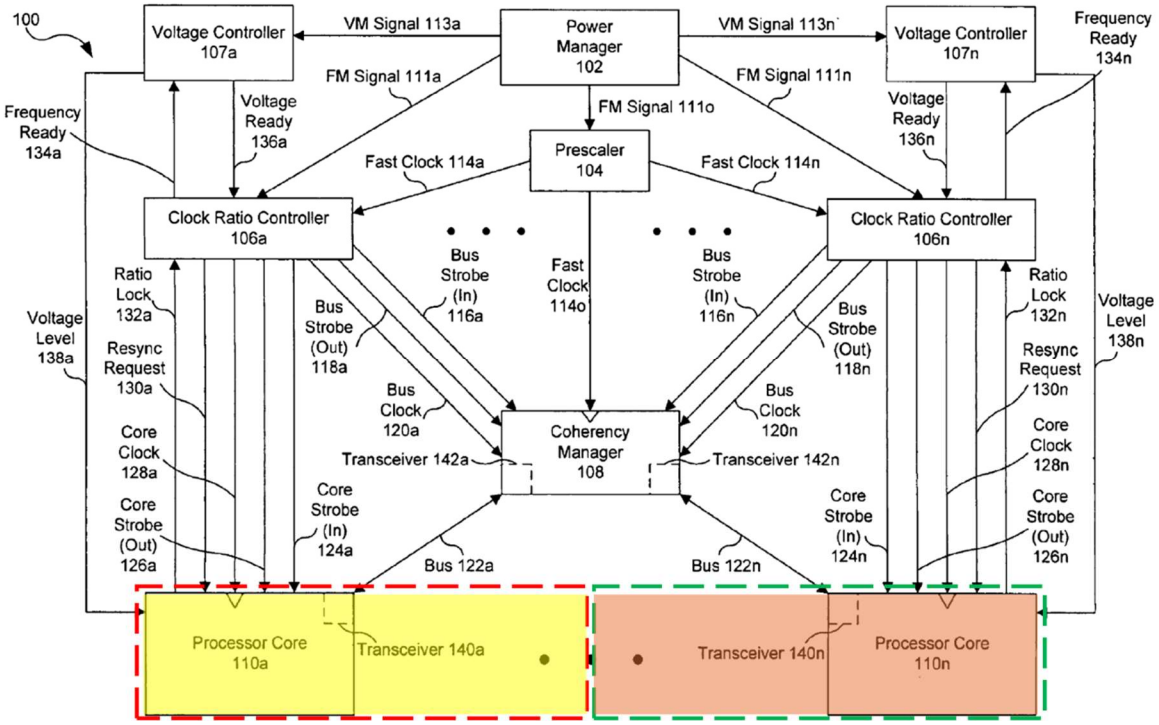


*FIG. 1A of Knoth, annotated*

179. In Allarey, the first set of cores (core 0 (308), core 1 (310), core 2 (312), and core 3 (314)) are located in the first region depicted below in red dotted box; the second set of cores (core 4 (316), core 5 (318), core 6 (320), and core 7 (322)) of Allarey are located in the second region depicted below in green dotted box. The first and second regions depicted below are completely separate in physical layout

and thus are non-overlapping. Accordingly, Allarey discloses this claim limitation.

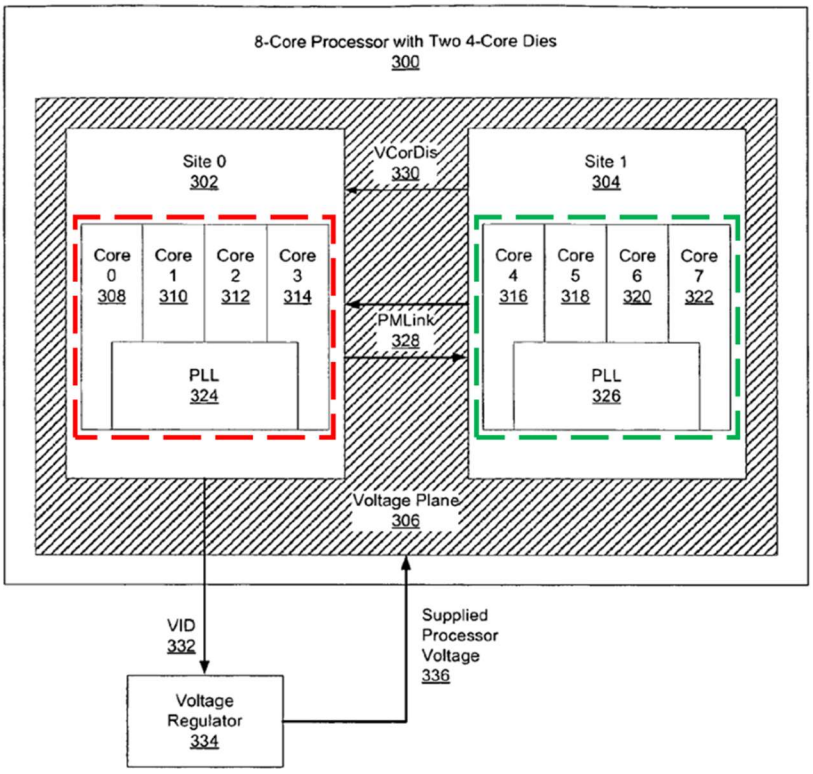


*FIG. 3 of Allarey, annotated*

## G. Dependent Claim 14

### 1. 14[pre] - The multi-core processor of claim 1,

180. As explained above for Claim 1, Knoth and Allarey, alone or in combination, teach each and every limitation of Claim 1, rendering it obvious.

### 2. 14[a] - wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a common region that is substantially central to the first set of processor cores and the second set of processor cores.

181. As explained above for Claim 5[a] in Ground 1, Knoth discloses "wherein the first set of processor cores and the second set of processor cores are

configured to receive one or more control signals from one or more control blocks." Specifically, Knoth discloses the control blocks in the form of power management unit 102 (highlighted in red) and clock ratio controllers 106a-n, as depicted below.

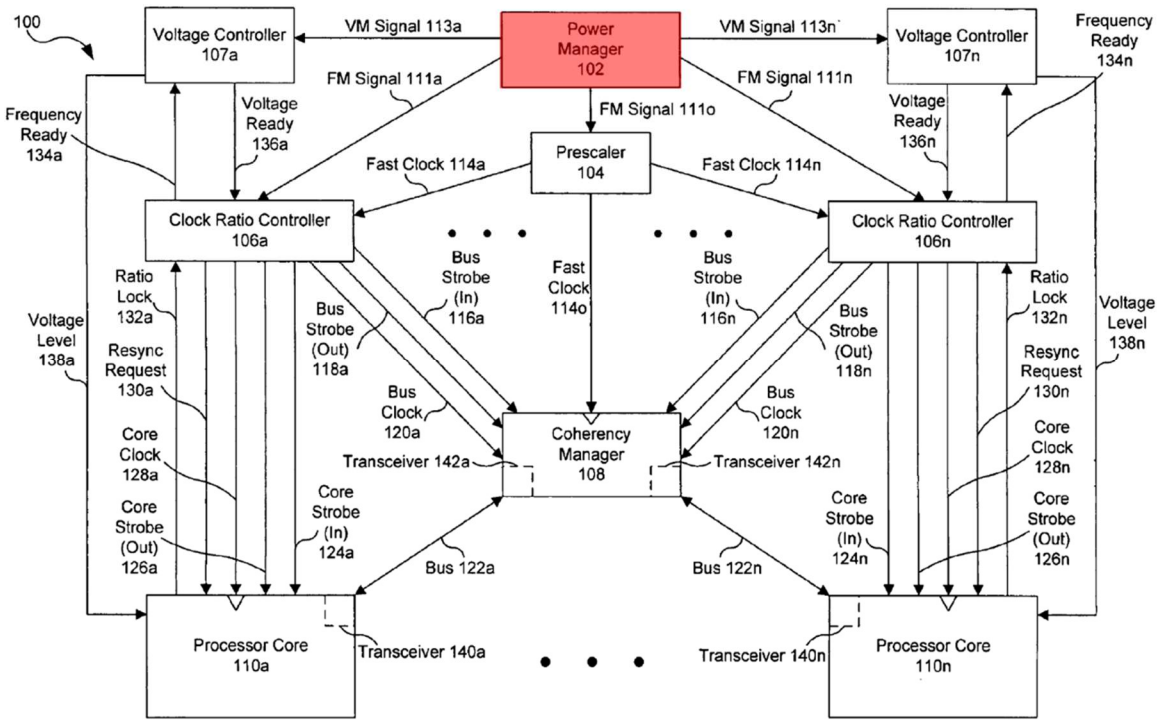


*FIG. 1A of Knoth, annotated*

182. As reviewed above for Claim 8 in Ground 1, the '339 Patent provides that a "region" is simply a spatial grouping of cores based on their physical location in the multi-core processor system. Power management unit 102 is positioned centrally relative to the first and second sets of cores Knoth (illustrated by processor core 110a and processor core 110n).

183. Furthermore, as explained above for Claim 5 in Ground 1, it was a known and obvious design option for a PHOSITA to place the control blocks of

Knoth in a common region central to the first and second sets of processor cores.

184. Allarey discloses the control blocks in the form of the logic within site 0 (102) and the control signals from voltage regulator 126. The logic within site 0 must be part of the circuitry within site 0. A PHOSITA would have understood that the logic within site 0 could be, for example, positioned towards the boundary of site 0 on the side of core 1, thus located substantially central to the first set of cores (illustrated by core 0 (108) and core 1 (110)) and the second set of cores (illustrated by core 0 (112) and core 1 (114)), as illustrated below by the red arrow. There must be circuitry that provides the control signals for the cores. This circuitry teaches the claimed one or more control signals.

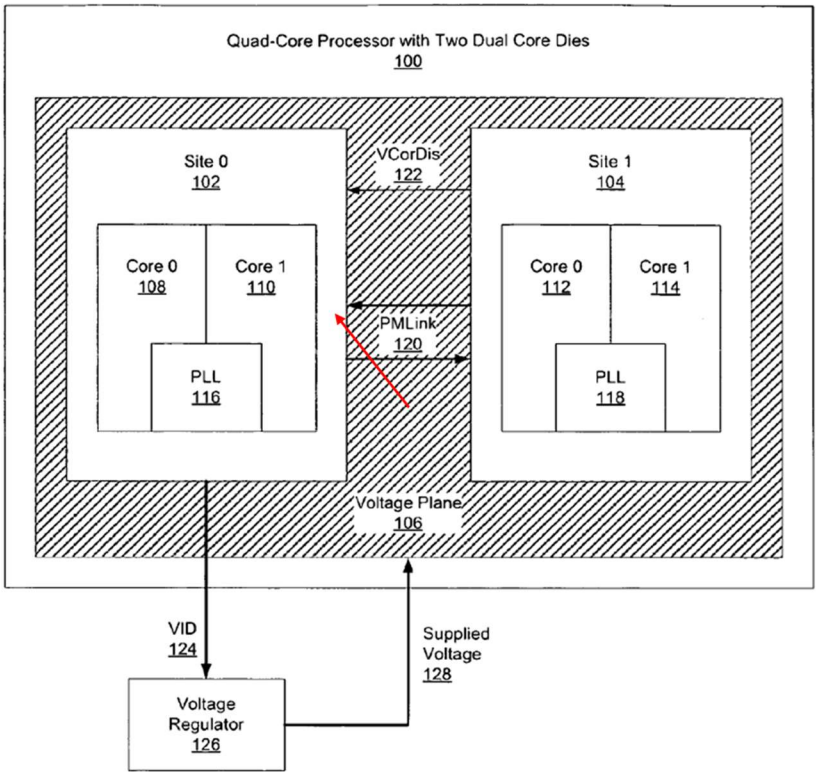

*FIG. 1 of Allarey, annotated*

185. Furthermore, as explained above for Claim 5 in Ground 1, it was a known and obvious design option for a PHOSITA to place the control blocks of Allarey in a common region central to the first and second sets of processor cores. Accordingly, Allarey discloses this claim limitation.

## H. Dependent Claim 21

### 1. 21[pre] - A multi-core processor, comprising:

186. To the extent the preamble is limiting, for the same reasons described for Claim 1[pre] in Ground 1, Knoth and Allarey, alone or in combination, disclose a multi-core processor.

### 2. 21[a1] - a first set of processor cores of the multi-core processor,

187. For the same reasons described for Claim 1[a1] in Ground 1, Knoth and Allarey, alone or in combination, disclose this limitation.

### 3. 21[a2] - wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage from a power control block and a first output clock signal from a first phase lock loop (PLL) having a first clock signal as input in a clock control block;

188. For the same reasons described for Claim 1[a2] in Ground 1, Knoth and Allarey, alone or in combination, disclose "wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage" and "a first output clock signal from a first phase lock loop (PLL) having a

first clock signal as input."

189. Knoth further discloses a power control block in the form of voltage controllers 107a-n (as shown in dashed blue block) and a clock control block in the form of clock ratio controllers 106a-n (as shown in dashed purple block). As discussed before, Knoth discloses that each processor core 110(a-n) receives its respective supply from voltage controller 107(a-n) and its respective output clock signals core clock signal 128(a-n) output from PLL 202 in clock ratio controller 106(a-n). Therefore, Knoth discloses this claim limitation.

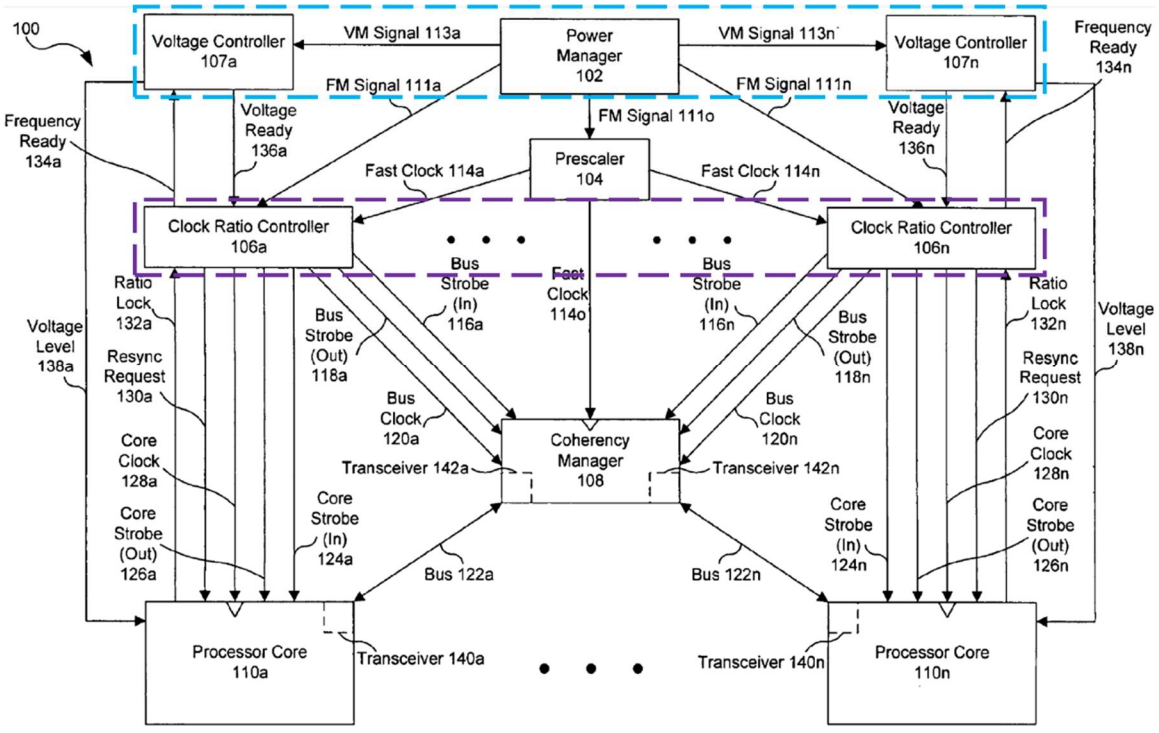


*FIG. 1A of Knoth, annotated*

190. Allarey further discloses a power control block in the form of voltage plane 106/306 (highlighted in red) and a clock control block in the form of PLL

circuits 116 and 118 (or 324 and 326) (as shown in dashed purple block). As discussed before, Allarey discloses that the first set of processor cores (for example, core 0 (108) and core 1 (110) in site 0 (102)) receives the first supply voltage from voltage plane 106, and the first output clock signal from PLL 116 in the PLL circuits that includes PLL 116 and 118. The second set of processor cores (for example, core 0 (112) and core 1 (114) in site 1 (104)) receives the second supply voltage from voltage plane 106, and the second output clock signal from PLL 118 in the PLL circuits that includes PLL 116 and 118. Therefore, Allarey discloses this claim limitation.

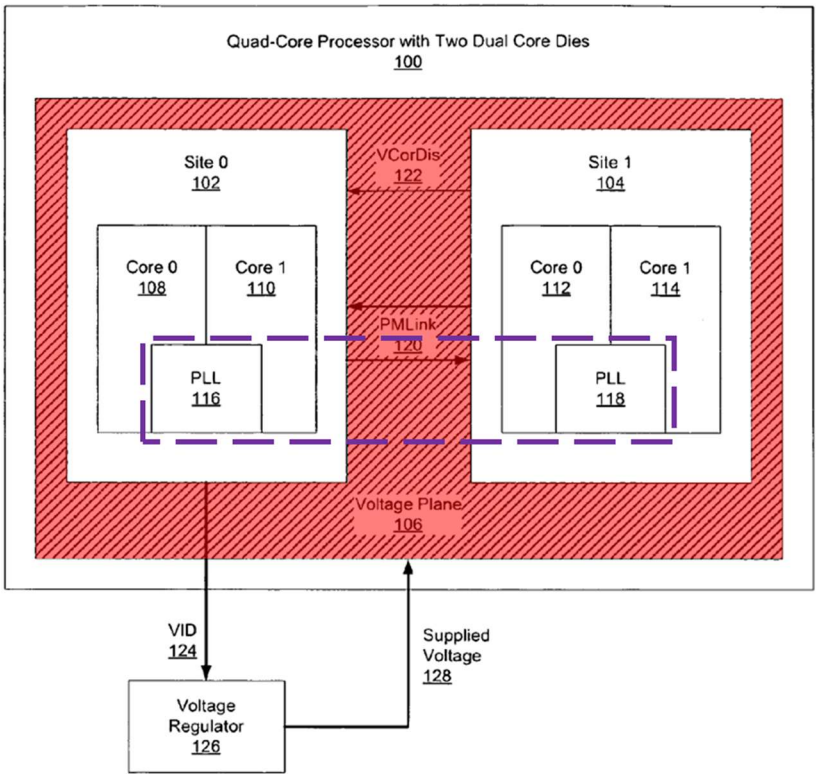


*FIG. 1 of Allarey, annotated*

**4. 21[b1] - a second set of processor cores of the multi-core processor,**

191. For the same reasons described for Claim 1[b1] in Ground 1 and as explained above for Claim 21[a1] in Ground 1, Knoth and Allarey, alone or in combination, disclose this claim limitation.

**5. 21[b2] - wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage from the power control block and a second output clock signal from a second PLL having a second clock signal as input in the clock control block,**

192. For the same reasons described for Claim 1[a2], [b2] in Ground 1 and as explained above for Claim 21[a2] in Ground 1, Knoth and Allarey, alone or in combination, disclose this claim limitation.

**6. 21[b3] - wherein the first supply voltage is independent from the second supply voltage, and**

193. For the same reasons described for Claim 1[b3] in Ground 1, Knoth and Allarey, alone or in combination, disclose this claim limitation.

**7. 21[b4] - the first clock signal is independent from the second clock signal; and**

194. For the same reasons described for Claim 1[b4] in Ground 1, Knoth and Allarey, alone or in combination, disclose this claim limitation.

**8. 21[c1] - an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores,**

195. For the same reasons described for Claim 1[c1] in Ground 1, Knoth and Allarey, alone or in combination, disclose this claim limitation.

**9. 21[c2] - wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores.**

196. For the same reasons described for Claim 1[c2] in Ground 1, Knoth and Allarey, alone or in combination, disclose this claim limitation.

## VII. GROUND 2: CLAIMS 2-4 ARE OBVIOUS UNDER §103 OVER KNOTH AND ALLAREY IN VIEW OF FLAUTNER

197. A PHOSITA would have understood that Knoth and Allarey, alone or in combination, further in view of Flautner teach each and every limitation of Claims 2-4.

### A. Motivation to Combine Knoth, Allarey and Flautner

198. A PHOSITA would have been motivated to combine Knoth and Allarey. *See* Section VI.A. A PHOSITA would have been further motivated to combine Knoth and Allarey with Flautner.

199. To improve the coordination and communication efficiency between the processor cores managed by Knoth's and Allarey's interface block, a PHOSITA would have been inclined to look beyond the teachings of these references to identify improved configurations. For example, a PHOSITA would have been inclined to seek references that cover specific components dedicated to the functions of clock synchronization and level shifting because the processor cores in Knoth and Allarey operate in different voltage domains with different voltage levels and at different frequencies. Knoth and Allarey expressly indicate the need for the interface block to

perform these functions. *See, e.g.*, Ex[1005] at [0002] ("communications and the exchange of data between various components are disrupted until a resynchronization occurs"), [0029]; Ex[1006] at 1:23-24 ("Asynchronous voltage changes during this time may disrupt a PLL lock process.").

200. Flautner, like Knoth and Allarey, discloses technology in the field of multi-core processor systems, specifically addressing the challenges of managing different operational states and power domains within such systems.

201. The coherency manager of Knoth couples the processor cores. Ex[1005] at [0033]. As a result, the processor cores can communicate with each other. This allows the processor cores to adjust the voltage levels and/or frequency of clock signals to optimize power consumption. Similarly, power management link of Allarey also couples site 0 and site 1 such that they can communicate with each other. Ex[1006] at 2:59-3:3. This allows for adjusting the voltage and/or clock and for the relocking process.

202. Flautner discloses the specific details of the components to handle clock synchronization and voltage level shifting between cores operating at different supply voltages and frequencies.

> FIG. 4 illustrates a fourth embodiment. ***This embodiment is similar to that of FIG. 1 except that in this case the first processor core 40 and the second processor core 42 are asymmetrically controlled by the clock speed controller 44. More particularly, this asymmetric control allows clocks of different speeds to be simultaneously supplied to respective ones of the first processor core 40 and the second processor***

> ***core 42. Thus, the first processor core 40 may be supplied with a relatively fast clock whilst the second processor core 42 is being supplied with a relatively slow clock.*** Accompanying these different speed clocks the integrated circuit upon which the first processor core 40 and the second processor core 42 are both formed may be split into multiple voltage domains with respective power controllers 46, 48. Thus, the second processor core 42 may have its clock speed reduced and its supply voltage lowered so as to reduce energy consumption by the second processor core 42 whilst the first processor core 40 maintains a high speed clock and a higher supply voltage needed for that high speed clock.
>
> When using different clock speeds and voltage levels in the first processor core 40 and the second processor core 42 it will be appreciated that ***a synchronisation module 50 and a voltage level shifter 52 is provided between the first processor core 40 and the second processor core 42 to deal with clock synchronisation issues and the different supply voltage levels (voltage signalling levels) between the two domains.***

Ex[1007] at [0047]-[0048].

203. Applying these teachings to Knoth and Allarey does not require substantial changes and would yield predictable results because such changes amount to a simple combination of known parts. In Knoth's architecture, the addition of Flautner's synchronization module and voltage level shifter would seamlessly integrate into the existing system of independent clock and voltage controllers, enhancing the system's capability to handle cores operating under diverse conditions. Similarly, in Allarey's structure, incorporating Flautner's components would address the potential voltage mismatches between sites and improve overall communication efficiency, ensuring that signals are properly synchronized and

translated between cores, regardless of their voltage and frequency differences.

204. For at least these reasons, a PHOSITA would have been motivated to combine Knoth and Allarey with Flautner.

### B. Dependent Claim 2

#### 1. 2[pre] - The multi-core processor of claim 1,

205. As explained above for Claim 1, Knoth and Allarey, alone or in combination, teach each and every limitation of Claim 1, rendering it obvious.

#### 2. 2[a] - wherein the interface block further comprises a first level shifter that is referenced to the second supply voltage and adapted to translate first logic levels associated with the first set of processor cores to second logic levels associated with the second set of processor cores for a first signal traveling from the first set of processor cores to the second set of processor cores.

206. Flautner teaches the use of the claimed level shifter. The '339 Patent discloses that the function of the level shifter is to "translate the signal levels such that each of the processor cores operates correctly (e.g., the processor cores properly interpret the voltages as valid logic levels even though processor cores are powered by different supply voltages)." Ex[1001] at 3:43-67. Similarly, Flautner discloses that the first processor core 40 and the second processor core 42 within separate voltage domains receiving different supply voltage levels. Ex[1007] at [0047]-[0048]. A voltage level shifter 52 is provided between the first processor core 40 and the second processor core 42 to "deal with … the different supply voltage levels (voltage signaling levels) between the two domains." *Id.* Voltage level shifter 52

thus translates the voltage of a signal from the cores in one voltage domain to the expected level associated with the cores in the other domain.

> FIG. 4 illustrates a fourth embodiment. ***This embodiment is similar to that of FIG. 1 except that in this case the first processor core 40 and the second processor core 42 are asymmetrically controlled by the clock speed controller 44. More particularly, this asymmetric control allows clocks of different speeds to be simultaneously supplied to respective ones of the first processor core 40 and the second processor core 42. Thus, the first processor core 40 may be supplied with a relatively fast clock whilst the second processor core 42 is being supplied with a relatively slow clock.*** Accompanying these different speed clocks the integrated circuit upon which the first processor core 40 and the second processor core 42 are both formed may be split into multiple voltage domains with respective power controllers 46, 48. Thus, the second processor core 42 may have its clock speed reduced and its supply voltage lowered so as to reduce energy consumption by the second processor core 42 whilst the first processor core 40 maintains a high speed clock and a higher supply voltage needed for that high speed clock.
>
> When using different clock speeds and voltage levels in the first processor core 40 and the second processor core 42 it will be appreciated that ***a synchronisation module 50 and a voltage level shifter 52 is provided between the first processor core 40 and the second processor core 42 to deal with clock synchronisation issues and the different supply voltage levels (voltage signalling levels) between the two domains.***

Ex[1007] at [0047]-[0048].

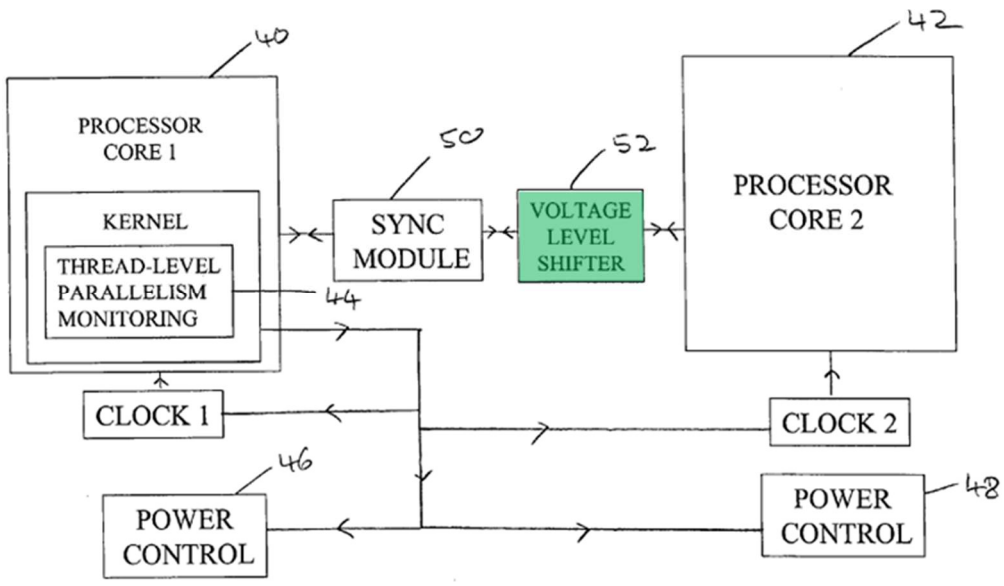


*Fig. 4 of Flautner, annotated*

207. A PHOSITA reading Flautner would have understood that translating voltage levels inherently involves translating the corresponding logic levels associated with the cores. This is because the supplied voltage levels correspond to logic levels within the cores and translating the voltage levels then results in translating the corresponding logic levels within the cores.

208. When a signal is transmitted between cores operating with different voltage levels, the signal's logic levels may not match the expected input levels of the receiving domain. Due to the discrepancy in the voltage levels, the logic levels between the two cores may not be consistent. For example, a logic level of '1' for a high voltage level will look different than a logic level of '0' for a low voltage level.

209. If a signal crosses from one domain to another without adjusting its voltage, it could be misinterpreted by the receiving core. Therefore, a voltage level

shifter must adjust the signal's voltage so that a logic '1' or '0' in one core remains a logic '1' or '0' in the other core, preserving the intended logic state despite different voltage requirements. The voltage level shifter acts as a translator between the two cores such that despite different voltage levels, the logic levels can still accurately be exchanged between the cores.

210. To the extent it is argued or found that voltage level shifter 52 has the shared functions between the processor cores in two voltage domains, it would have been obvious to a PHOSITA that voltage level shifter 52 could have separate level shifter modules for the core(s) in each voltage domain because level shifters are standard, routine components and Flautner already has it in the multi-core processor system, and because adding a known level shifter in the processor layout of Knoth or Allarey would require this simple modification. For each voltage domain, a separate level shifter module is used. In this way, every voltage domain is afforded a level shifter module for shifting the logic levels in that voltage domain. A PHOSITA would have also been motivated to provide dedicated level shifters for each domain to further improve precision of voltage translation across voltage domains. Additional level shifters across voltage domains would improve the accuracy of the voltage translation.

### C. Dependent Claim 3

#### 1. 3[pre] - The multi-core processor or claim 1,

211. As explained above for Claim 1, Knoth and Allarey, alone or in combination, teach each and every limitation of Claim 1, rendering it obvious.

#### 2. 3[a] - wherein the interface block further comprises a second level shifter that is referenced to the first supply voltage and adapted to translate second logic levels associated with the second set of processor cores to first logic levels associated with the first set of processor cores for a second signal traveling from the second set of processor cores to the first set of processor cores.

212. For the same reasons described for Claim 2[a] in Ground 2, Flautner discloses this claim limitation. Accordingly, Knoth and Allarey, further in view of Flautner, thus teach each and every limitation of Claim 3, rendering it obvious.

### D. Dependent Claim 4

#### 1. 4[pre] - The multi-core processor or claim 1,

213. As explained above for Claim 1 in Ground 1, Knoth and Allarey, alone or in combination, teach each and every limitation of Claim 1, rendering it obvious.

#### 2. 4[a] - wherein the interface block further comprises a synchronizer configured to synchronize the first clock signal and the second clock signal for communication between one or more processor cores of the first set of processor cores and one or more processor cores of the second set of processor cores.

214. Flautner teaches the use of a synchronizer. The ’339 Patent discloses that the function of the level shifter is to “synchronize the clock signal[s] for the

communication between the processor core[s].” Ex[1001] at 4:64-5:8. Specifically, Flautner discloses that a synchronisation module 50 is provided between the first processor core 40 and the second processor core 42. Ex[1007] at [0048]. The first processor core 40 and the second processor core 42 are “asymmetrically controlled by the clock speed controller 44” that supplies “clocks of different speeds … to respective ones of the first processor core 40 and the second processor core 42.” *Id.* at [0047]. The synchronisation module “deal[s] with clock synchronisation issues” between the first processor core 40 and the second processor core 42 running at different clock speeds.

> FIG. 4 illustrates a fourth embodiment. ***This embodiment is similar to that of FIG. 1 except that in this case the first processor core 40 and the second processor core 42 are asymmetrically controlled by the clock speed controller 44. More particularly, this asymmetric control allows clocks of different speeds to be simultaneously supplied to respective ones of the first processor core 40 and the second processor core 42. Thus, the first processor core 40 may be supplied with a relatively fast clock whilst the second processor core 42 is being supplied with a relatively slow clock.*** Accompanying these different speed clocks the integrated circuit upon which the first processor core 40 and the second processor core 42 are both formed may be split into multiple voltage domains with respective power controllers 46, 48. Thus, the second processor core 42 may have its clock speed reduced and its supply voltage lowered so as to reduce energy consumption by the second processor core 42 whilst the first processor core 40 maintains a high speed clock and a higher supply voltage needed for that high speed clock.
>
> When using different clock speeds and voltage levels in the first processor core 40 and the second processor core 42 it will be appreciated that ***a synchronisation module 50*** and a voltage level shifter 52 ***is provided between the first processor core 40 and the***

> ***second processor core 42 to deal with clock synchronisation issues*** and the different supply voltage levels (voltage signalling levels) between the two domains.

Ex[1007] at [0047]-[0048].

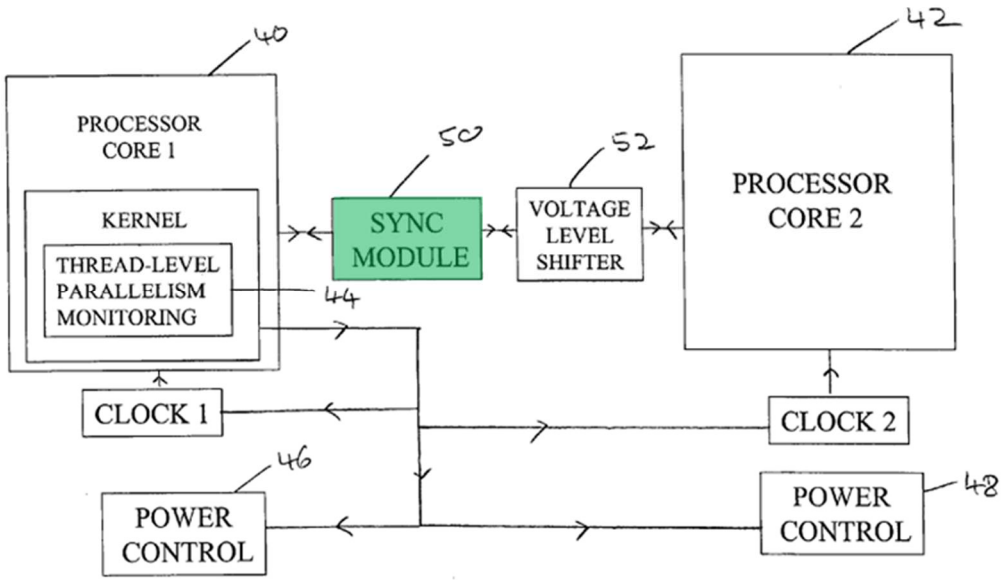


*Fig. 4 of Flautner, annotated*

215. A PHOSITA would have understood that to facilitate communication between the first and second cores of Knoth and Allarey, the clock signals of the two cores would necessarily be required to be synchronized. Data can only be exchanged between cores when their clock signals are aligned. When two cores operate at different clock speeds, the timing of signal transfers between the cores will not align. When the timing of signal transfers are not aligned, then data will not be able to be transferred from one core to the other. Synchronization of the timing of the clock signals is necessary to make sure the communication from one core is recognized and correctly interpreted by the other core. Otherwise, there will be no

communication between the cores. The synchronisation module 50 of Flautner thus synchronizes the clock signals of the first and second cores for communication between the first and second cores, as would have been required in the systems of Knoth and Allarey. Flautner's synchronisation module allows the core clock speeds to be synchronized such that communication between the cores is open. Knoth and Allarey, further in view of Flautner, thus teach each and every limitation of Claim 4, rendering it obvious.

## VIII. GROUND 3: CLAIM 6 IS OBVIOUS UNDER §103 OVER KNOTH AND ALLAREY IN VIEW OF WOLFE, AND FURTHER IN VIEW OF KUMAR

216. A PHOSITA would have understood that Knoth and Allarey, alone or in combination, in view of Wolfe and/or Kumar teaches or suggests each and every limitation of Claim 6 and therefore renders it obvious.

### A. Motivation to Combine Knoth, Allarey and Wolfe

217. A PHOSITA would have been motivated to combine Knoth and Allarey. *See* Section VI.A. A PHOSITA would have been further motivated to combine Knoth and Allarey with Wolfe.

218. A PHOSITA would have recognized that Knoth and Allarey focus more on functional mechanisms for dynamic voltage and frequency regulation in a multi-core processor system. It would have been in the best interest of a PHOSITA to explore complementary configurations that optimize the physical layout of the multi-

core processor system to support these mechanisms. Wolfe discloses a grid organization and relevant physical layout considerations for a multi-core processor having independent control mechanisms for dynamic management of voltage and frequency levels. See, e.g., Ex[1008] at [0014]. This structured layout would have prompted a PHOSITA to consider combining the functional mechanisms disclosed by Knoth and Allarey with the physical layout strategy of Wolfe to achieve a multi-core processor system that maximizes both operational efficiency and performance.

**B. Motivation to Combine Knoth, Allarey and Kumar**

219. A PHOSITA would have been further motivated to combine Knoth and Allarey with Kumar. To improve the power management and inter-core voltage regulation of Knoth's and Allarey's multi-core processor systems, a PHOSITA would have further been inclined to look beyond the teachings of these references to identify improved configurations as part of the normal course of his/her own research. A PHOSITA would have been inclined to seek references that cover specific mechanisms for inter-core voltage regulation in a multi-core processor.

220. As discussed above for Claim 5 in Ground 1, Knoth and Allarey disclose control blocks for managing voltage and frequency independently for processor cores. *See, e.g.*, Ex[1005] at [0025], Ex[1006] at 2:65-3:19. Kumar, like Knoth and Allarey, discloses a power management system in the context of integrated circuits. As reviewed in more detail above, Kumar discloses additional

details for dynamically managing the respective voltage supplies within predetermined voltage relationships to optimize the overall power distribution and performance. Ex[1009] at [0054]-[0060]. A PHOSITA would have been motivated to combine Knoth and Allarey further with Kumar because Kumar's teaching could improve Knoth's and/or Allarey's system by providing a more granular control of voltage levels across different sets of cores to improve voltage stability and power efficiency.

221. Applying these teachings from Kumar to Knoth and Allarey does not require substantial changes and would yield predictable results because such changes amount to a simple combination of known parts.

### C. Dependent Claim 6

#### 1. 6[pre] - The multi-core processor of claim 5,

222. As explained above for Claim 5, Knoth and Allarey, alone or in combination, teach each and every limitation of Claim 5, rendering it obvious.

#### 2. 6[a] - wherein the first set of processor cores is adjacent to the second set of processor cores, and

223. Knoth and Allarey, alone or in combination, disclose "wherein the first set of processor cores is adjacent to the second set of processor cores."

224. Knoth discloses that processor cores 110a-n are placed adjacent to each other, as depicted below.

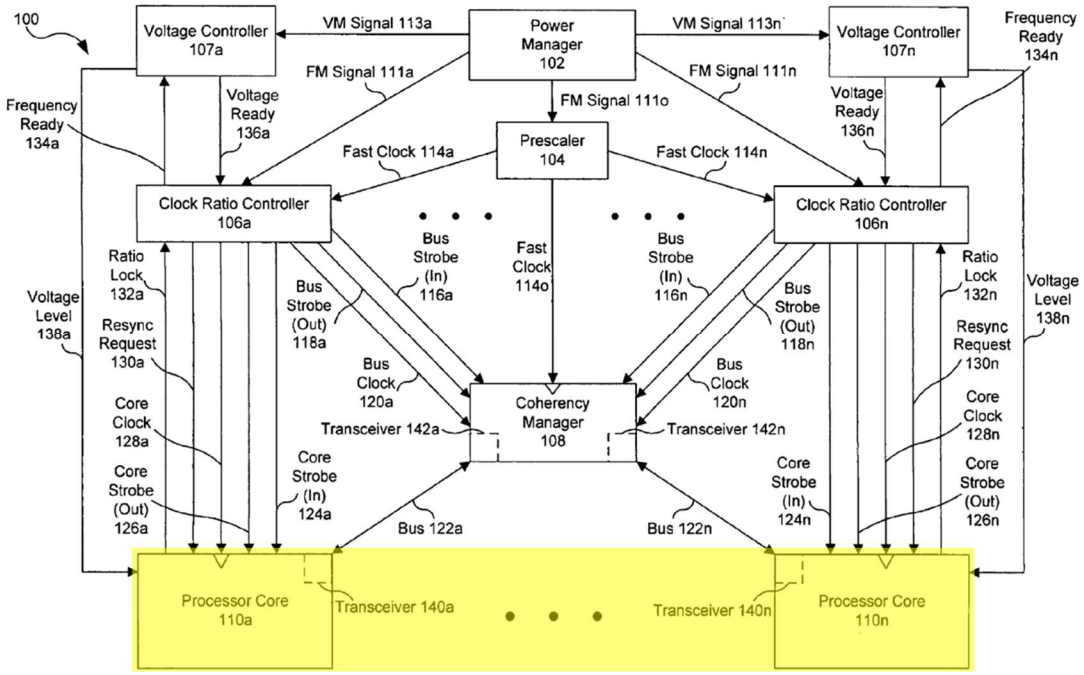


*Fig. 1A of Knoth, annotated*

225. Allarey likewise discloses that the cores in sites 0 and 1 are adjacent as depicted below.

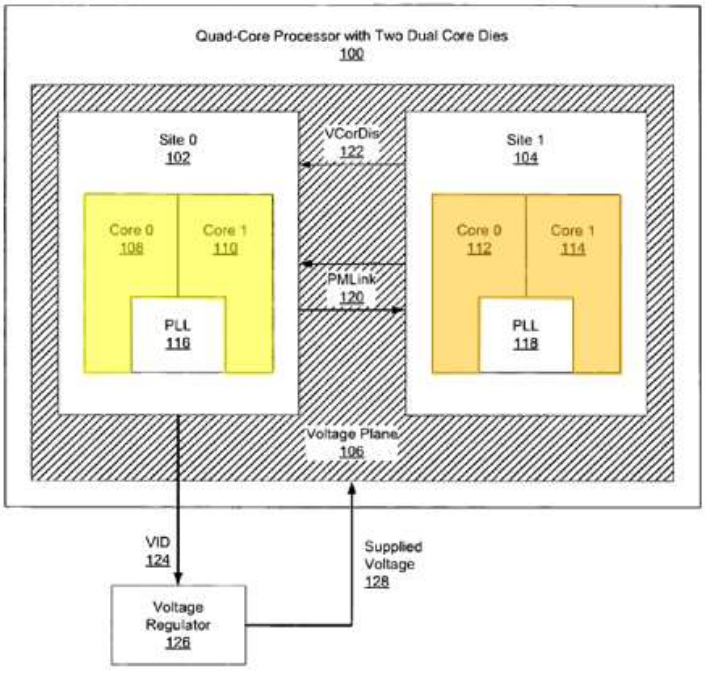


*FIG. 1 of Allarey, annotated*

226. Wolfe also discloses this claim limitation. Wolfe discloses a grid organization of processor cores, and grouping cores based on geometric mapping.

> FIG. 1A is an illustration of a processor core 100 in a multi-core processor 102 adapted to operate based on a selected voltage level and a clock frequency, arranged in accordance to at least some embodiments of the present disclosure. ***The multi-core processor 100 may include multiple processor cores arranged in rows and columns in a 2-dimensional array.*** The processor cores may be interconnected by a communication network using switches, which is illustrated in FIG. 3. A processor core may be supplied a certain voltage level (e.g., a selected voltage level 124) from a voltage control circuit (e.g., a voltage control circuit 110) and a clock signal at a certain clock frequency (e.g., a selected clock frequency 126) from a clock control circuit (e.g., a clock control circuit 112).

Ex[1008] at [0014].

> When the computing task 104 is divided, a geometric mapping signal 106 and a computing task signal 108 may be generated by a host processor, or one or more designated processor cores. The computing task signal 108 may convey information associated with one or more of the divided computing tasks, and the geometric mapping signal 106 may convey information associated with the geometric mapping between the divided computing task and a processor core. The information associated with the computing tasks may include pixel data, reference frame data, motion estimation data, motion compensation data, or other data that forms the basis of the computing task. ***The information associated with the geometric mapping may include the assignment of a particular divided computing task to one or more processor cores of the multi-core processor.***

*Id.* at [0016].

> Processing for the method 200 may begin at block 202, ***"Determine First Workload for First Computing Task Geometrically Mapped to First Processor Core,"*** where the method 200 may be arranged to determine the first workload of the first processor core in the multi-core processor. ***The first workload is for the first processor core to perform a first computing task associated with a first image area and a first***

***geometric mapping between the first computing task and the first processor core.***

*Id.* at [0022].

227. The first set of cores of the multi-core processor could be, for example, located in a first region corresponding to a row of the two-dimensional array (highlighted in yellow), and the second set of cores in a second region corresponding to an adjacent row of the array (highlighted in orange).

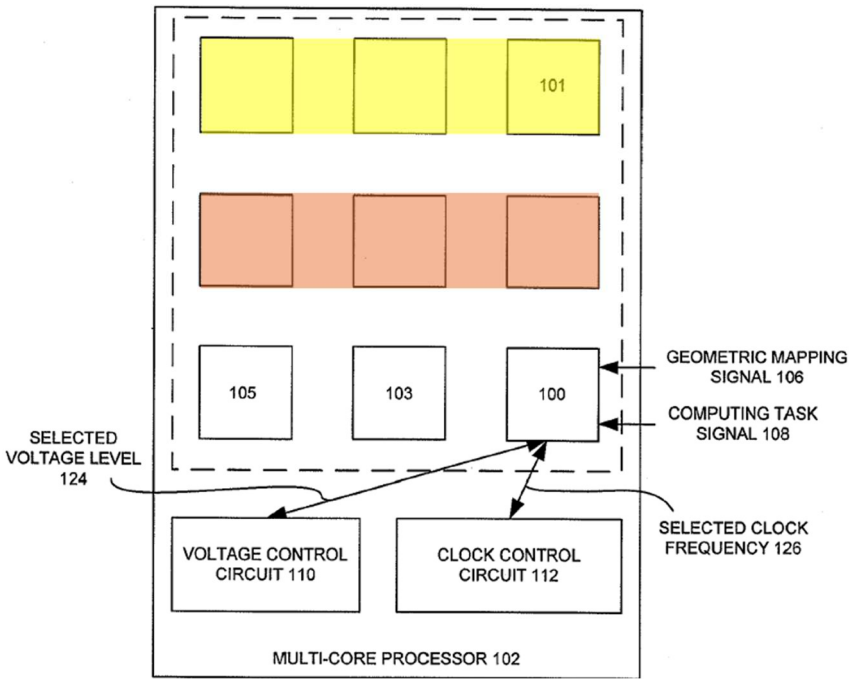


*Figure 1 of Wolfe, modified*

**3. 6[b] - the one or more control blocks are configured to select the first supply voltage and the second supply voltage to maintain a differential relationship between the first supply voltage and the second supply voltage.**

228. Kumar discloses "one or more control blocks" that "select the first supply voltage and the second supply voltage to maintain a differential relationship

between the first supply voltage and the second supply voltage." Specifically, Kumar discloses managing and maintaining a differential relationship between the voltages of two sets of components using a controller. Kumar discloses a controller 540 that controls "generat[ing] a voltage V1 across one or more components 130 and … a voltage V2 across one or more components 530, wherein voltage V1 is to satisfy one or more predetermined relationships with voltage V2." *Id.* at [0054]. "Controller 540 … identif[ies] whether a voltage V1 across one or more components 130 and a voltage V2 across one or more components 530 satisfy one or more predetermined relationships," including "whether the absolute value of the difference between voltage V1 and voltage V2 is less than, or less than or equal to, a predetermined amount or a predetermined percentage of either voltage V1 or voltage V2." Ex[1009] at [0058]. "If voltage V1 and voltage V2 do not satisfy one or more predetermined relationships," controller 540 makes adjustments to "help voltage V1 and voltage V2 satisfy one or more predetermined relationships." *Id.* at [0059]-[0060].

> Controller 540 for one embodiment may be coupled to control programmable current sources 546 and 548 to measure a resistance Rpackage of package 120 based at least in part on a reference resistance of package 120, the amount of current I1 drawn through one or more components 130, and the amount of current 12 drawn through one or more components 530. ***Controller 540 for one embodiment may control programmable current source 546 to generate a voltage V1 across one or more components 130 and may control programmable current source 548 to generate a voltage V2 across one or more components 530, wherein voltage V1 is to satisfy one or more predetermined relationships with voltage V2.***

*Id.* at [0054].

> For block 706, ***controller 540 may identify whether a voltage V1 across one or more components 130 and a voltage V2 across one or more components 530 satisfy one or more predetermined relationships. Controller 540 for one embodiment may identify for block 706 whether the absolute value of the difference between voltage V1 and voltage V2 is less than, or less than or equal to, a predetermined amount or a predetermined percentage of either voltage V1 or voltage V2.*** Controller 540 for one embodiment may identify for block 706 whether voltage V1 and voltage V2 are substantially equal. For one embodiment where one or more components 130 and one or more components 530 are coupled to a common supply voltage node 501, controller 540 for one embodiment may identify for block 706 whether measured voltages from nodes 502 and 503 satisfy one or more predetermined relationships.

*Id.* at [0058].

> ***If voltage V1 and voltage V2 do not satisfy one or more predetermined relationships for block 706, controller 540 for block 708 may control programmable current source 546 to adjust current I1 through one or more components 130 and/or control programmable current source 548 to adjust current I2 through one or more components 530. Controller 540 for one embodiment for block 708 may adjust current I1 and/or current I2 in any suitable manner to help voltage V1 and voltage V2 satisfy one or more predetermined relationships for block 706.*** Controller 540 may repeat operations for blocks 706 and 708 until voltage V1 and voltage V2 satisfy one or more predetermined relationships for block 706.
>
> If voltage V1 and voltage V2 do satisfy one or more predetermined relationships for block 706, controller 540 for block 710 may measure a ratio between a resistance R1 of one or more components 130 and a resistance R2 of one or more components 530 based at least in part on the amount of current I1 and the amount of current I2 to help measure a resistance Rpackage of package 120. Controller 540 for one embodiment may measure such a ratio in accordance with the following equation.
>
> $R1/R2 = (V1*I2)/(V2*I1)$

*Id.* at [0059]-[0060].

229. The architecture disclosed by Kumar is applicable to a multi-core processor system. *See, e.g., id.* at [0073].

## IX. GROUND 4: CLAIM 11 IS OBVIOUS UNDER §103 OVER KNOTH AND ALLAREY IN VIEW OF WOLFE

230. A PHOSITA would have understood that Knoth in view of Allarey and further in view of Wolfe teaches or suggests each and every limitation of Claim 11 and therefore renders it obvious.

### A. 11[pre] - The multi-core processor or claim 8,

231. As explained above for Claim 8 in Ground 1, Knoth and Allarey, alone or in combination, teach each and every limitation of Claim 8, rendering it obvious.

### B. 11[a] - wherein the first region corresponds to a first row of the multi-core processor, and wherein the second region corresponds to a second row of the multi-core processor.

232. For the same reasons described for Claim 6[a] in Ground 3, a PHOSITA reading Wolfe would have understood that the first set of cores of the multi-core processor thus could be located in a first region corresponding to a first row of the two-dimensional array (highlighted in yellow), and the second set of cores in a second region corresponding to a second row of the array (highlighted in orange).

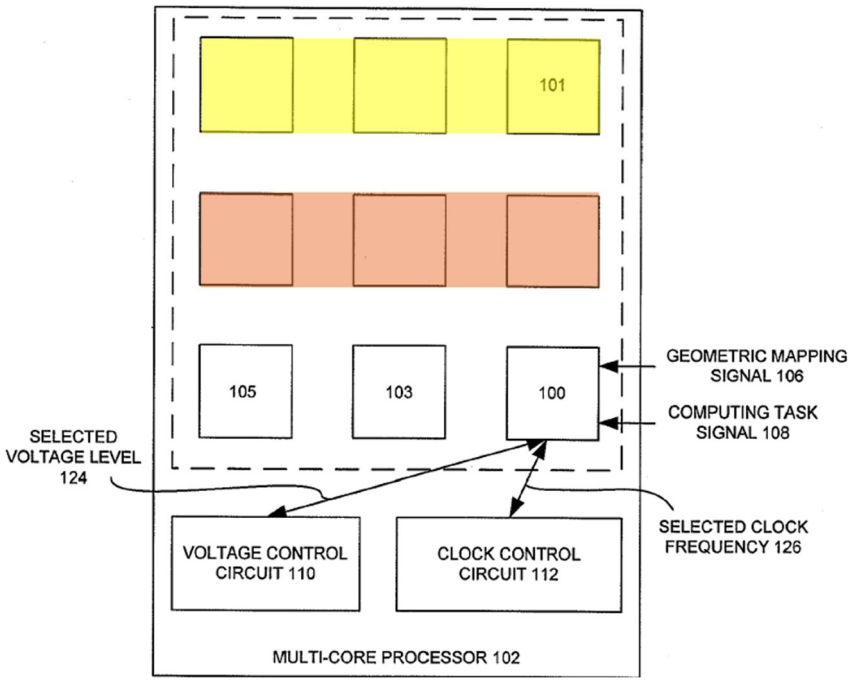


*Figure 1A of Wolfe, modified*

233. Accordingly, Knoth and/or Allarey in view of Wolfe renders Claim 11 obvious.

## X. GROUND 5: CLAIMS 1-3, 5, 8-10, 14 AND 21 ARE OBVIOUS UNDER §103 OVER NAFFZIGER IN VIEW OF ALLAREY

234. Naffziger and Allarey, alone or in combination, teach each and every limitation of Claims 1-3, 5, 8-10, 14 and 21 and therefore renders these claims obvious.

### A. Motivation to Combine Naffziger and Allarey

235. A PHOSITA would have been motivated to combine Naffziger and Allarey and would have had a reasonable expectation of success in doing so, because they relate to the same well-known technologies.

236. First, both Naffziger and Allarey are directed to the same field of multi-

core processors, and they each address similar problems and propose similar solutions for managing voltage and frequency scaling of multi-core processors, similar to those disclosed in the '339 Patent. For example, Naffziger seeks to propose solutions for multi-core processor systems' control over power consumption that include changing operational states with different power consumptions by adjusting the supply voltage, clock frequency and other parameters. Ex[1010] at [0001]-[0012]. Like Naffziger, Allarey seeks to propose solutions for optimizing multi-core processor systems' power conservation by dynamically modifying the voltage supplied to and the frequency of the processor. *See* Ex[1006] at 1:6-24.

237. Naffziger is authored by Advanced Micro Devices engineers, and Allarey is authored by Intel engineers. A PHOSITA would have looked to publications by such leading companies in designing multi-core processing systems, such as AMD and Intel, and considered the different techniques disclosed in these references for optimizing the power and/or performance of cores.

238. A PHOSITA would, at a minimum, have found the combination obvious to try because it combines well-known techniques that are inter-related—Naffziger deals with dynamic power and performance adjustments, and Allarey studies stabilizing a supplied voltage during a clock signal frequency locking process. A PHOSITA would have found synergy in combining dynamic power adjustments from Naffziger with voltage stabilization from Allarey to improve

overall processor performance and stability. Therefore, a PHOSITA would have found that Naffziger and Allarey provide complimentary solutions and would have been motivated to combine them to benefit from their respective teachings.

239. Finally, a PHOSITA would have had a considerable expectation of success when combining these teachings because the combination would involve a mere substitution of one known element (e.g., Naffziger's core-level control) with another (Allarey's site-level control), applying a known technique to a known system ready for improvement, and/or use of known techniques to improve performance in a similar multi-core environment. Therefore, the teachings and considerations of Allarey would have allowed a PHOSITA to improve on Naffziger's systems effortlessly (and vice versa). For at least these reasons, a PHOSITA would have been motivated to seek and combine Naffziger and Allarey.

### B. Independent Claim 1

#### 1. 1[pre] - A multi-core processor, comprising:

240. To the extent that the preamble is limiting, Naffziger teaches this subject matter. Naffziger discloses a multi-core processor (e.g., dual core processor [500], having a first core 105A and a second core 105B). *See, e.g.*, Ex[1010] at FIG. 5 (processor cores highlighted in yellow).

> A processor is also disclosed. ***The processor includes one or more processor cores*** and a power control unit. The power control unit is configured to, for a time interval of operation in a first operational state, determine an amount of power consumed during by the one or more

> processor cores, calculate a power error based on the amount of power consumed in the time interval, add the power error for the time interval to a power error term from a previous time interval, and compare the power error term for the time interval to a first error threshold. The power control unit is further configured to cause the processor to exit the first operational state and enter a second operational state when the power error term for the time interval is outside a range defined at least in part by the first error threshold. If the power error term for the time interval is within the range defined at least in part by the first error threshold, the power control unit is configured to cause the processor to remain in the first operational state.

Ex[1010] at [0008].

> Another embodiment is contemplated wherein ***a multi-core processor includes multiple power control units, each of which is associated with a corresponding one of a plurality of processor cores.*** In such an embodiment, each power control unit may separately control the states of operation of its corresponding core. The power control unit for each core may also monitor power consumption of the processor as a whole, which may be used, along with one or more power error terms, as a basis for determining the operational state of its corresponding processor core.

*Id.* at [0055].

> Turning now to FIG. 5, a block diagram of one embodiment of a processor having a power control unit and a plurality of processor cores is shown. In the embodiment shown, ***processor 500 is a dual core processor, having a first core 105A and a second core 105B***. Processor 500 also includes a north bridge unit 110 and a power control unit 150, both of which are coupled to both processor cores 105A and 105B. Power control unit 150 is configured to monitor the power of both processor cores 105A and 105B, and configured to alternate operation of these cores between at least a high performance state and a low power state. As noted above, additional, intermediate operational states may be implemented. Power control unit 150 may accomplish the switching between operational states of processor cores 105A and 105B by changing the frequency of clock signals provided to the cores, changing supply voltages provided to the cores, or both.

*Id.* at [0053].

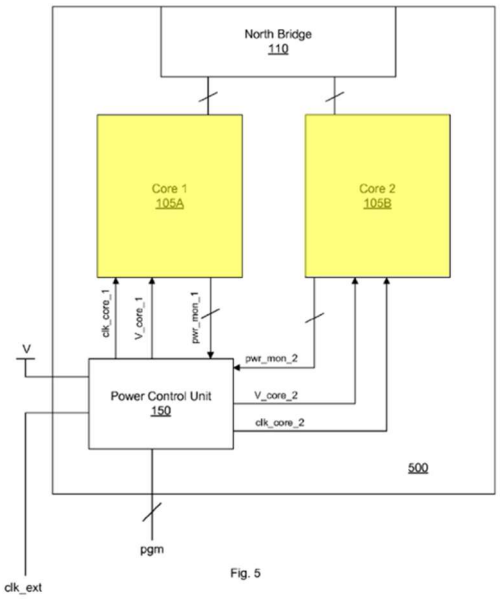


*Fig. 5 of Naffziger, annotated*

241. As explained above for Claim 1[Preamble] in Ground 1, Allarey likewise teaches a multi-core processor.

**2. 1[a1] - a first set of processor cores of the multi-core processor,**

242. Naffziger and Allarey, alone or in combination, teach "a first set of processor cores of the multi-core processor" and "a second set of processor cores of the multi-core processor."

243. Naffziger discloses that the multi-core processor has a first core 105A and a second core 105B "the operational states [of which are] controlled independently of one another." Ex[1010] at [0054]. For the first and second cores 105A and 105B, Naffziger discloses a power control unit 150 with separate,

independent supply voltage and clock signal control planes, or alternatively, multiple power control units each associated with a corresponding process core, that regulates their core supply voltage and core clock frequency.

> Turning now to FIG. 5, a block diagram of one embodiment of a processor having a power control unit and a plurality of processor cores is shown. In the embodiment shown, processor 500 is a dual core processor, having a first core 105A and a second core 105B. Processor 500 also includes a north bridge unit 110 and a power control unit 150, both of which are coupled to both processor cores 105A and 105B. ***Power control unit 150 is configured to monitor the power of both processor cores 105A and 105B, and configured to alternate operation of these cores between at least a high performance state and a low power state.*** As noted above, additional, intermediate operational states may be implemented. ***Power control unit 150 may accomplish the switching between operational states of processor cores 105A and 105B by changing the frequency of clock signals provided to the cores, changing supply voltages provided to the cores, or both.***
>
> In some embodiments, the operational state of cores 105A and 105B are in conjunction with each other. That is, both processor cores 105A and 105B are operated in the same state. However, ***embodiments are possible and contemplated wherein the operational states of processor cores 105A and 105B may be controlled independently of one another.*** For example, in such an embodiment, processor core 105A may operate in the high performance state, while processor core 105B operates in the low power state. ***Such a configuration may require that cores 105A and 105B have separate, independent core supply voltage planes and/or are configured to receive separate, independently controlled core clock signals.*** In addition, embodiments are possible and contemplated wherein power control unit 150 is configured to monitor the processing workloads of each processor core 105A and 105B, and may change operational states of the cores by changing their respective processing workloads.

*Id.* at [0053]-[0054].

> In the embodiment shown, power control unit 150 is configured to regulate and control power consumption by processor 100. More particularly, power control unit 150 is configured to control the amount of power consumed by processor core 105 in order to keep the average power consumption, over time, within a predetermined limit. In this particular embodiment, power control unit is configured to control power consumption by core 105 by controlling a voltage, controlling a clock frequency, or both. Power control unit is coupled to receive voltage V from a source external to processor 100. Similarly, power control unit is coupled to receive an external clock signal, clk_ext, from a source external to processor 100. In turn, power control unit is coupled to provide a core supply voltage, V_core, and a core clock frequency, clk_core, to processor core 105. ***Thus, power consumption by processor core 105 may be varied by varying the level of the core supply voltage or the frequency of the clock signal.***

*Id.* at [0022].

> Power control unit 150 in the embodiment shown is configured to monitor the power consumed by processor core 105 and to regulate power consumption. In the embodiment shown, one or more signals may be conveyed from processor core 105 to power control unit 150 via a power monitoring bus, pwr_mon. The signals conveyed may include indications of various operating parameters of processor core 105. Such parameters may include, but are not limited to, processor core temperature, current consumption, processing load, bus activity (i.e. activity between North Bridge 110 and processor core 105) and so forth. Based on one or more of these parameters, ***power control unit 150 may adjust the core supply voltage or the frequency of the core clock signal in order to control the power consumption of core 105, with the goal of balancing power consumption with processor performance.***

*Id.* at [0023].

> In one embodiment, an algorithm for determining an operating state of processor 100 includes comparing the power error term with at least one error threshold. ***If processor 100 is operating in a high performance state, the error threshold may represent a power error deficit, and if this deficit is exceeded, a change of operational states may occur. If processor 100 is operating in a low power state, the***

> ***error threshold may represent a power error surplus, which, if exceeded, may also cause a change in operational states.*** If processor 100 is operating in an intermediate state (e.g., one or one or more states between a high performance state and a low power state), comparisons of a value stored in register 224 to two different error thresholds may be performed by control logic 220. If one of these values is exceeded (i.e. the power error term is outside the range defined by the error thresholds), processor 100 may enter another operational state. If the power error term is within this range (e.g., neither a power error surplus nor a power error deficit threshold is exceeded), operation of processor 100 can continue in the present state. Various embodiments of algorithms for determining a processor operating state will be discussed in further detail below.

*Id.* at [0030].

> Another embodiment is contemplated wherein a multi-core processor includes multiple power control units, each of which is associated with a corresponding one of a plurality of processor cores. In such an embodiment, ***each power control unit may separately control the states of operation of its corresponding core.*** The power control unit for each core may also monitor power consumption of the processor as a whole, which may be used, along with one or more power error terms, as a basis for determining the operational state of its corresponding processor core.

*Id.* at [0055].

244. To the extent it is argued or found that the first and second cores 105A and 105B are individual cores, it would have been obvious to a PHOSITA to extend the disclosed architecture of Naffziger to include two distinct sets of cores, where each set of cores is monitored and controlled separately. For example, instead of managing individual cores, the system could treat core 105A and additional cores associated with it as a first set, while core 105B and its associated cores could form a second set. Each set would have been monitored and controlled independently,

allowing the system to fine-tune power consumption based on the specific demands of each set. As reviewed above for Claim 1[a1] in Ground 1, it would have been obvious for a PHOSITA to extend Naffziger's architecture to include two sets of cores, which was a routine modification that would have been easily implemented.

245. Accordingly, Naffziger discloses and/or renders obvious "a first set of processor cores of the multi-core processor" and "a second set of processor cores of the multi-core processor."

246. Furthermore, a PHOSITA would have been motivated to modify Naffziger to have a first and second sets of cores, rather individual first and second cores, based on the teachings as disclosed in Allarey. As explained above for Claim 1[a1] in Ground 1, Allarey discloses a first set of processor cores in the form of a multi-core die in site 0. Ex[1006] at 2:41-58, 5:17-33.

247. Therefore, Naffziger and Allarey, alone or in combination, disclose this claim limitation.

**3. 1[a2] - wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage and a first output clock signal of a first phase lock loop (PLL) having a first clock signal as input;**

248. Naffziger discloses that the first core 105A receives a first voltage supply "V_core_1" (shown by the blue line) and a first output clock signal "clk_core_1" (shown by the purple line) of a phase locked loop (PLL) of the power control unit 150 (highlighted in purple) having a first clock signal "clk_ext" (shown

by the dark blue line) as input.

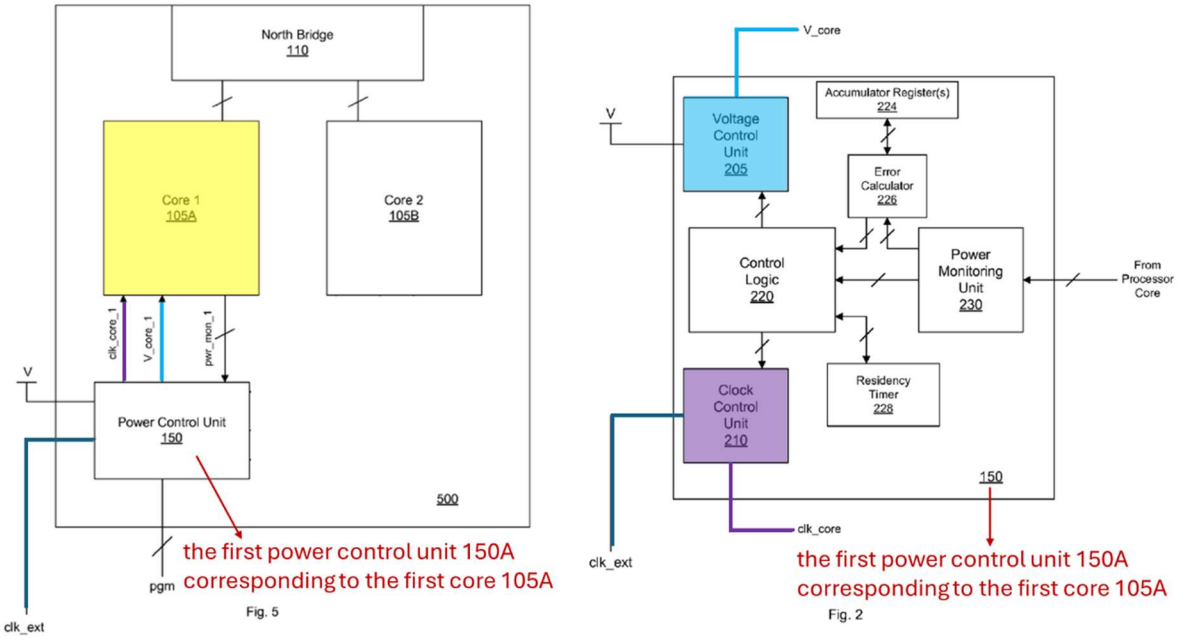


*Figs. 2 and 5 of Naffziger, annotated*

In the embodiment shown, power control unit 150 is configured to regulate and control power consumption by processor 100. More particularly, power control unit 150 is configured to control the amount of power consumed by processor core 105 in order to keep the average power consumption, over time, within a predetermined limit. In this particular embodiment, power control unit is configured to control power consumption by core 105 by controlling a voltage, controlling a clock frequency, or both. Power control unit is coupled to receive voltage V from a source external to processor 100. Similarly, power control unit is coupled to receive an external clock signal, clk_ext, from a source external to processor 100. In turn, ***power control unit is coupled to provide a core supply voltage, V_core, and a core clock frequency, clk_core, to processor core 105***. Thus, power consumption by processor core 105 may be varied by varying the level of the core supply voltage or the frequency of the clock signal.

Turning now to FIG. 5, a block diagram of one embodiment of a processor having a power control unit and a plurality of processor cores is shown. In the embodiment shown, processor 500 is a dual core processor, having a first core 105A and a second core 105B. ***Processor***

> ***500 also includes a north bridge unit 110 and a power control unit 150, both of which are coupled to both processor cores 105A and 105B. Power control unit 150 is configured to monitor the power of both processor cores 105A and 105B, and configured to alternate operation of these cores between at least a high performance state and a low power state***. As noted above, additional, intermediate operational states may be implemented. ***Power control unit 150 may accomplish the switching between operational states of processor cores 105A and 105B by changing the frequency of clock signals provided to the cores, changing supply voltages provided to the cores, or both.***

*Id.* at [0022], [0053].

249. Further, Naffziger discloses that the cores receive corresponding supply voltages from the respective power control units.

> ***Another embodiment is contemplated wherein a multi-core processor includes multiple power control units, each of which is associated with a corresponding one of a plurality of processor cores. In such an embodiment, each power control unit may separately control the states of operation of its corresponding core.*** The power control unit for each core may also monitor power consumption of the processor as a whole, which may be used, along with one or more power error terms, as a basis for determining the operational state of its corresponding processor core.

Ex[1010] at [0055].

250. A PHOSITA reading Naffziger would have understood that the multi-core processor 500 includes more than one power control unit 150 each associated with a corresponding core—a first power control unit 150 corresponding to the first core 105A (for ease of discussion, referred to as the "first power control unit 150A") and a second power control unit 150 corresponding to the second core 105B (for ease of discussion, referred to as the "second power control unit 150B"), as

illustrated above in annotated Figs 2 and 5.

251. Naffziger discloses dynamically "adjust[ing] the core supply voltage or the frequency of the core clock signal." Ex[1010] at [0023]. Specifically, Naffziger discloses that the first power control unit 150A includes the voltage control unit 205 that "us[es] adjustable voltage regulator circuitry, level shifter circuity, or any other suitable circuitry configured to vary a supply voltage," including "V_core_1." *Id.* at [0032], Fig. 2 (reproduced above, annotated). Accordingly, Naffziger discloses that its first set of cores dynamically receives a first supply voltage ("V_core_1") that varies as adjusted by circuitry of the voltage control unit 205 of the first power control unit 150A.

> Power control unit 150 in the embodiment shown is configured to monitor the power consumed by processor core 105 and to regulate power consumption. In the embodiment shown, one or more signals may be conveyed from processor core 105 to power control unit 150 via a power monitoring bus, pwr_mon. The signals conveyed may include indications of various operating parameters of processor core 105. Such parameters may include, but are not limited to, processor core temperature, current consumption, processing load, bus activity (i.e. activity between North Bridge 110 and processor core 105) and so forth. Based on one or more of these parameters, ***power control unit 150 may adjust the core supply voltage or the frequency of the core clock signal in order to control the power consumption of core 105, with the goal of balancing power consumption with processor performance.***

*Id.* at [0023].

> Control logic 220 in the embodiment shown is coupled to both a voltage control unit 205 and a clock control unit 210. ***Control of the operating state of processor 100 may be accomplished by control logic 220***

> ***through manipulating one or both of a core supply voltage (V_core) or the frequency of a core clock signal (clk_core).*** Control logic 220 may alter the operating state of processor 100 between the low power state by directing voltage control unit 205 to change the core supply voltage provided to processor core 105. ***Voltage control unit 205 may be implemented using adjustable voltage regulator circuitry, level shifter circuitry, or any other suitable circuitry configured to vary a supply voltage.*** An external voltage V may be received by voltage control unit 205 from a power supply via a motherboard to which processor 100 may be coupled. In one embodiment, control logic 220 may provide signals to voltage control unit 205 that cause processor 100 to enter the low power state by reducing the core supply voltage, and may cause processor 100 to enter the high performance state by increasing the core supply voltage.

*Id.* at [0032].

252. Each power control unit 150 including the first power control unit 150A further includes the clock control unit 210 that "us[es] a phase locked loop (PLL) … for adjusting the frequency of a clock signal" and "is configured to vary the frequency of the core clock signal" to dynamically provide a "core clock signal, clk_core, … to processor core 105," including "clk_core_1." *Id.* at [0033], Figs. 2 and 5 (reproduced above, annotated). Accordingly, Naffziger discloses that its first set of cores dynamically receives a first output clock signal ("clk_core_1") of a first PLL (the PLL used by the clock control unit 210 of the first power control unit 150A, shaded in purple).

> ***Clock control unit 210 may be implemented using a phase locked loop (PLL)***, a delay locked loop (DLL), or other type of circuitry that can be used ***for adjusting the frequency of a clock signal.*** In the embodiment shown, an external clock signal, clk_ext, is provided at a first frequency to clock control unit 210. ***A core clock signal, clk_core, is provided by***

> ***clock control unit 210 to processor core 105. Clock control unit 210 is configured to vary the frequency of the core clock signal based on signals received from control logic 220.*** In one embodiment, control logic 220 may cause processor 100 to enter the low power state by directing clock control unit 210 to reduce the frequency of the core clock signal, and may cause the processor 100 to enter the high performance state by directing clock control unit to increase the frequency of the core clock signal.

*Id.* at [0033].

253. The PLL of clock control unit 210 of the first power control unit 150A receives as an input "an external clock signal, clk_ext." *Id.* at [0033], Figs. 2 and 5 (reproduced above, annotated); *see also id.* at [0022] (each power control unit 150 "is coupled to receive an external clock signal, clk_ext.").

> Clock control unit 210 may be implemented using a phase locked loop (PLL), a delay locked loop (DLL), or other type of circuitry that can be used for adjusting the frequency of a clock signal. ***In the embodiment shown, an external clock signal, clk_ext, is provided at a first frequency to clock control unit 210.*** A core clock signal, clk_core, is provided by clock control unit 210 to processor core 105. Clock control unit 210 is configured to vary the frequency of the core clock signal based on signals received from control logic 220. In one embodiment, control logic 220 may cause processor 100 to enter the low power state by directing clock control unit 210 to reduce the frequency of the core clock signal, and may cause the processor 100 to enter the high performance state by directing clock control unit to increase the frequency of the core clock signal.

*Id.* at [0033].

> In the embodiment shown, power control unit 150 is configured to regulate and control power consumption by processor 100. More particularly, power control unit 150 is configured to control the amount of power consumed by processor core 105 in order to keep the average power consumption, over time, within a predetermined limit. In this particular embodiment, power control unit is configured to control

> power consumption by core 105 by controlling a voltage, controlling a clock frequency, or both. Power control unit is coupled to receive voltage V from a source external to processor 100. ***Similarly, power control unit is coupled to receive an external clock signal, clk_ext, from a source external to processor 100.*** In turn, power control unit is coupled to provide a core supply voltage, V_core, and a core clock frequency, clk_core, to processor core 105. Thus, power consumption by processor core 105 may be varied by varying the level of the core supply voltage or the frequency of the clock signal.

*Id.* at [0022].

254. As explained above for Claim 1[a2] in Ground 1, Allarey also discloses this claim limitation.

**4. 1[b1] - a second set of processor cores of the multi-core processor,**

255. As explained above for Claim 1[a1] in Ground 5, Naffziger discloses "a second set of processor cores of the multi-core processor." Furthermore, to the extent it is argued or found that Naffziger does not disclose this limitation, a PHOSITA would have been motivated to modify the architecture of Naffziger to include a first and second sets of cores in view of Allarey. Therefore, Naffziger and Allarey, alone or in combination, disclose "a second set of processor cores of the multi-core processor."

**5. 1[b2] - wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage and a second output clock signal of a second PLL having a second clock signal as input,**

256. Naffziger discloses "wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage and

a second output clock signal of a second PLL having a second clock signal as input."

257. As depicted in Figs. 2 and 5, Naffziger discloses that the second core 105B receives a second voltage supply "V_core_2" (shown by the blue line) and a second output clock signal "clk_core_2" (shown by the purple line) of a phase locked loop (PLL) of the second power control unit 150B (highlighted in purple) having a second clock signal "clk_ext" (shown by the dark blue line) as input.

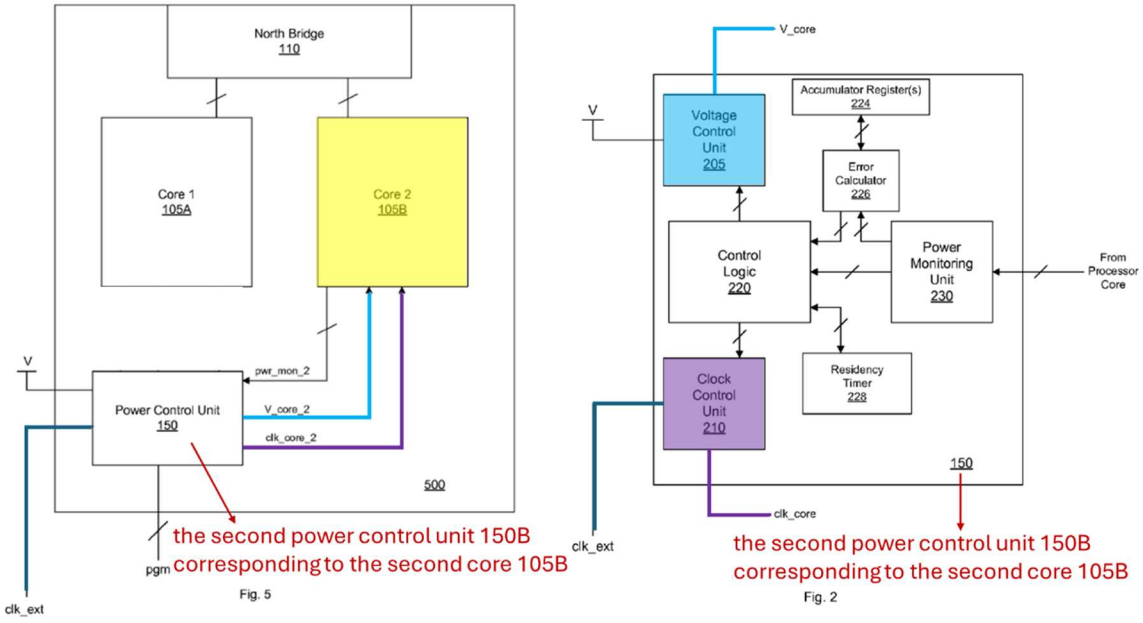


*Figs. 2 and 5 of Naffziger, annotated*

258. The second core 105B of Naffziger receives the second voltage supply "V_core_2" and the second output clock signal "clk_core_2" from the power control unit 150 (including the second power control unit 150B). As explained above for Claim 1[a2] in Ground 5, Naffziger discloses the power control unit includes circuitry to "vary a supply voltage"—including "V_core_2." Ex[1010] at 0032], Fig.

2 (reproduced above, annotated). Accordingly, Naffziger discloses that its second set of cores dynamically receives a second supply voltage ("V_core_2").

> Control logic 220 in the embodiment shown is coupled to both a voltage control unit 205 and a clock control unit 210. Control of the operating state of processor 100 may be accomplished by control logic 220 through manipulating one or both of a core supply voltage (V_core) or the frequency of a core clock signal (clk_core). Control logic 220 may alter the operating state of processor 100 between the low power state by directing voltage control unit 205 to change the core supply voltage provided to processor core 105. ***Voltage control unit 205 may be implemented using adjustable voltage regulator circuitry, level shifter circuitry, or any other suitable circuitry configured to vary a supply voltage.*** An external voltage V may be received by voltage control unit 205 from a power supply via a motherboard to which processor 100 may be coupled. In one embodiment, control logic 220 may provide signals to voltage control unit 205 that cause processor 100 to enter the low power state by reducing the core supply voltage, and may cause processor 100 to enter the high performance state by increasing the core supply voltage.

*Id.* at [0032].

259. The power control unit 150 (including the second power control unit 150B further includes the clock control unit 210 that "us[es] a phase locked loop (PLL) … for adjusting the frequency of a clock signal" and "is configured to vary the frequency of the core clock signal"—including "clk_core_2." *Id.* at [0033], Fig. 2 (reproduced above, annotated). Accordingly, Naffziger discloses that its second set of cores dynamically receives a second output clock signal ("clk_core_2") of a second PLL (*i.e.*, the PLL used by the clock control unit 210 of the second power control unit 150B).

> ***Clock control unit 210 may be implemented using a phase locked loop (PLL)***, a delay locked loop (DLL), or other type of circuitry that can be used ***for adjusting the frequency of a clock signal***. ***In the embodiment shown, an external clock signal, clk_ext, is provided at a first frequency to clock control unit 210.*** A core clock signal, clk_core, is provided by clock control unit 210 to processor core 105. Clock control unit 210 is configured to vary the frequency of the core clock signal based on signals received from control logic 220. In one embodiment, control logic 220 may cause processor 100 to enter the low power state by directing clock control unit 210 to reduce the frequency of the core clock signal, and may cause the processor 100 to enter the high performance state by directing clock control unit to increase the frequency of the core clock signal.

*Id.* at [0033].

260. The PLL of clock control unit 210 of the second power control unit 150A receives as an input "an external clock signal, clk_ext." *Id.* at [0033], Figs. 2 and 5 (reproduced above, annotated); *see also id.* at [0022] (each power control unit 150 "is coupled to receive an external clock signal, clk_ext.").

> In the embodiment shown, power control unit 150 is configured to regulate and control power consumption by processor 100. More particularly, power control unit 150 is configured to control the amount of power consumed by processor core 105 in order to keep the average power consumption, over time, within a predetermined limit. In this particular embodiment, power control unit is configured to control power consumption by core 105 by controlling a voltage, controlling a clock frequency, or both. Power control unit is coupled to receive voltage V from a source external to processor 100. ***Similarly, power control unit is coupled to receive an external clock signal, clk_ext, from a source external to processor 100.*** In turn, power control unit is coupled to provide a core supply voltage, V_core, and a core clock frequency, clk_core, to processor core 105. Thus, power consumption by processor core 105 may be varied by varying the level of the core supply voltage or the frequency of the clock signal.

*Id.* at [0022].

261. As explained above for Claim 1[b2] in Ground 1, Allarey also discloses this claim limitation.

**6. 1[b3] - wherein the first supply voltage is independent from the second supply voltage, and**

262. Naffziger discloses "wherein the first supply voltage is independent from the second supply voltage." Naffziger discloses that the first and second sets of cores receive independent voltage supplies and clock signals. Naffziger discloses that "the operational states of processor cores 105A and 105B may be controlled independently of one another." Ex[1010] at [0054]. "Such a configuration may require that cores 105A and 105B have separate, independent core supply voltage planes and/or are configured to receive separate, independently controlled core clock signals." *Id.*

> In some embodiments, the operational state of cores 105A and 105B are in conjunction with each other. That is, both processor cores 105A and 105B are operated in the same state. However, ***embodiments are possible and contemplated wherein the operational states of processor cores 105A and 105B may be controlled independently of one another.*** For example, in such an embodiment, processor core **105**A may operate in the high performance state, while processor core **105**B operates in the low power state. ***Such a configuration may require that cores 105A and 105B have separate, independent core supply voltage planes and/or are configured to receive separate, independently controlled core clock signals.*** In addition, embodiments are possible and contemplated wherein power control unit 150 is configured to monitor the processing workloads of each processor core 105A and 105B, and may change operational states of the cores by changing their respective processing workloads.

*Id.* at [0054].

263. As explained above for Claim 1[b3] in Ground 1, Allarey also discloses this claim limitation.

**7. 1[b4] - the first clock signal is independent from the second clock signal; and**

264. Naffziger discloses "the first clock signal is independent from the second clock signal." As discussed above for Claim 1[a2] and [b2]-[b3] in Ground 5, Naffziger discloses that the multi-core processor system includes separate, independent power control units associated with the cores—for example, the first power control unit 150A corresponding to the first core 105A and the second power control unit 150B corresponding to the second core 105B—and each power control unit receives as an input an external clock signal. A PHOSITA would have understood Naffziger to include independent first and second power control units receiving the first and second clock signals that are independent from each other.

265. As explained above for Claim 1[b4] in Ground 1, Allarey likewise discloses this claim limitation.

**8. 1[c1] - an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores,**

266. Naffziger discloses "an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores."

267. Naffziger discloses that the power control units are associated with their

respective corresponding cores, and further discloses that "[t]he power control unit for each core … also monitor power consumption of the processor as a whole … as a basis for determining the operational state of its corresponding core." Ex[1010] at [0055]. A PHOSITA would have understood that the power control units 150A and 150B would have been coupled to both the first and second sets of cores to monitor the multi-core processor as a whole.

> Another embodiment is contemplated wherein a multi-core processor includes multiple power control units, each of which is associated with a corresponding one of a plurality of processor cores. In such an embodiment, each power control unit may separately control the states of operation of its corresponding core. ***The power control unit for each core may also monitor power consumption of the processor as a whole, which may be used, along with one or more power error terms, as a basis for determining the operational state of its corresponding processor core.***

*Id.* at [0055].

268. Alternatively, Naffziger discloses a "power control unit 150 [that] is configured to monitor the processing workloads of each processor core 105A and 105B" such that "processor cores 105A and 105B may be controlled independently of one another" with "separate, independent core supply voltage planes" for the voltage supply and clock signals, as depicted in Fig. 5 (below). *See, e.g.*, Ex[1010] at [0054]. Naffziger thus also discloses a power control unit (highlighted in green) coupled to both sets of processor cores (highlighted in yellow) with separate, independent planes dedicated to each set.

> In some embodiments, the operational state of cores 105A and 105B are in conjunction with each other. That is, both processor cores 105A and 105B are operated in the same state. However, ***embodiments are possible and contemplated wherein the operational states of processor cores 105A and 105B may be controlled independently of one another. For example, in such an embodiment, processor core 105A may operate in the high performance state, while processor core 105B operates in the low power state. Such a configuration may require that cores 105A and 105B have separate, independent core supply voltage planes and/or are configured to receive separate, independently controlled core clock signals.*** In addition, embodiments are possible and contemplated wherein ***power control unit 150 is configured to monitor the processing workloads of each processor core 105A and 105B***, and may change operational states of the cores by changing their respective processing workloads.

*Id.* at [0054].

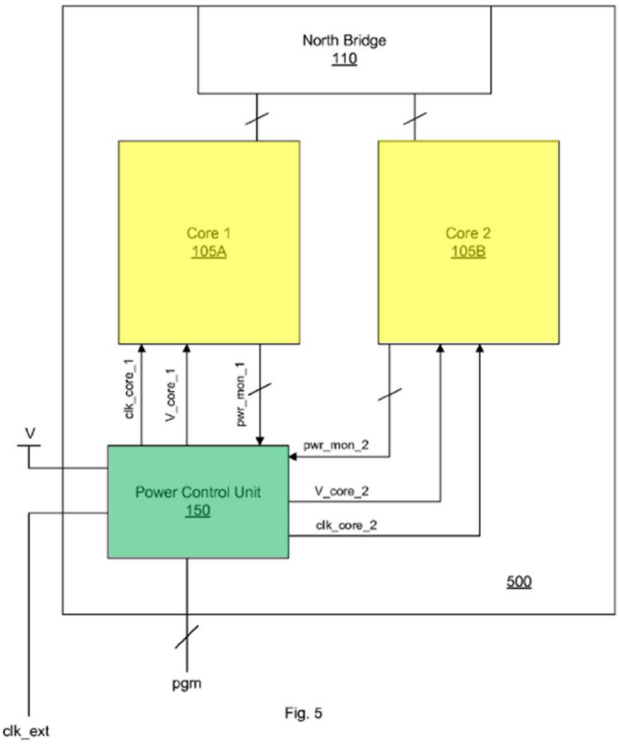


*Fig. 5 of Naffziger, annotated*

269. Naffziger teaches that the planes of the power control unit each controls

a core separately and independently with independent operational state with independent voltage supply and clock signals. Ex[1010] at [0054]. A PHOSITA reading Naffziger would have understood that for each core, the corresponding plane controls its voltage and frequency independently of the other core. A PHOSITA thus would have understood that for each core, its corresponding plane contains the entire architecture of at least an independent voltage control unit 205 to regulate its power supply and an independent clock control unit 210 to control clock signals.

270. The dual-plane power control unit 150 (shown in green box) would have appeared as depicted below, with separate planes for controlling the first core 105A and second core 105B independently.

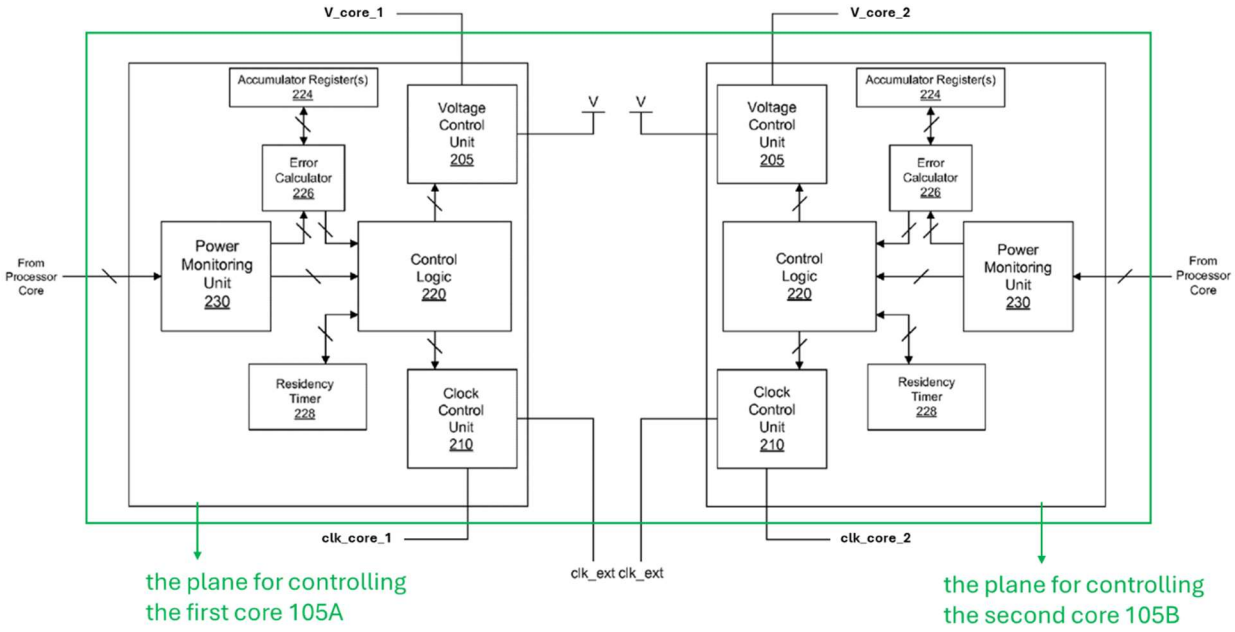


*Fig. 2 of Naffziger, modified to show separate planes for controlling the first core 105A and second core 105B independently*

271. As explained above for Claim 1[b4] in Ground 1, Allarey likewise discloses this claim limitation.

**9. 1[c2] - wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores.**

272. The interface block of Naffziger (the power control units or the dual-plane power control unit 150) is "configured to facilitate communication between the first set of processor cores and the second set of processor cores."

273. Naffziger discloses that the "[p]ower control unit 150 is configured to monitor the power of both processor cores 105A and 105B, and configured to alternate operation of these cores." Ex[1010] at [0053]-[0054]. Furthermore, the power control unit(s) "monitor power consumption of the processor as a whole … as a basis for determining the operational state of [each] processor core." Ex[1010] at [0055]. A PHOSITA would have understood that the power information of the processor as a whole including that of the second core could be communicated to the first core for determining the operational state of the first core, and vice versa.

> Turning now to FIG. 5, a block diagram of one embodiment of a processor having a power control unit and a plurality of processor cores is shown. In the embodiment shown, processor 500 is a dual core processor, having a first core 105A and a second core 105B. Processor 500 also includes a north bridge unit 110 and a power control unit 150, both of which are coupled to both processor cores 105A and 105B. ***Power control unit 150 is configured to monitor the power of both processor cores 105A and 105B, and configured to alternate operation of these cores between at least a high performance state and a low power state.*** As noted above, additional, intermediate operational

> states may be implemented. Power control unit 150 may accomplish the switching between operational states of processor cores 105A and 105B by changing the frequency of clock signals provided to the cores, changing supply voltages provided to the cores, or both.
>
> In some embodiments, the operational state of cores **105**A and **105**B are in conjunction with each other. That is, both processor cores **105**A and **105**B are operated in the same state. However, ***embodiments are possible and contemplated wherein the operational states of processor cores 105A and 105B may be controlled independently of one another.*** For example, in such an embodiment, processor core **105**A may operate in the high performance state, while processor core **105**B operates in the low power state. ***Such a configuration may require that cores 105A and 105B have separate, independent core supply voltage planes and/or are configured to receive separate, independently controlled core clock signals.*** In addition, embodiments are possible and contemplated wherein power control unit **150** is configured to monitor the processing workloads of each processor core **105**A and **105**B, and may change operational states of the cores by changing their respective processing workloads.
>
> Another embodiment is contemplated wherein a multi-core processor includes multiple power control units, each of which is associated with a corresponding one of a plurality of processor cores. In such an embodiment, each power control unit may separately control the states of operation of its corresponding core. ***The power control unit for each core may also monitor power consumption of the processor as a whole, which may be used, along with one or more power error terms, as a basis for determining the operational state of its corresponding processor core.***

*Id.* at [0053]-[0055].

274. To the extent it is argued or found that Naffziger does not expressly disclose that the interface block facilitates communications between the cores, as explained above for Claim 1[b4] in Ground 1, Allarey discloses an interface block in the form of the power management link (PMLink) 120/328 that "communicatively

couples" the first set of cores (the cores within site 0) and the second set of cores (the cores within site 1) and "transmit[s] data back and forth between" the first and second sets of cores.

> In many embodiments, ***a power management link (PMLink) 120 communicatively couples site 0 and site 1.*** The specific details of the PMLink 120 and its interface to each site can comprise one of many different link (i.e. interconnect, bus) forms. Generally, ***the PMLink 120 is capable of transmitting data back and forth between site 0 (102) and site 1 (104).*** In many embodiments, there is a master site and a slave site with respect to controlling the voltage level supplied to the voltage plane 106. In many embodiments, site 0 (102) is capable of controlling the voltage level supplied to the voltage plane 106. The voltage control process may be referred to as voltage correction.

Ex[1006] at 2:59-3:3.

275. Accordingly, Naffziger and Allarey, alone or in combination, disclose this limitation, as well as the remaining limitations of Claim 1 as discussed above, rendering Claim 1 obvious.

### C. Dependent Claim 2

#### 1. 2[pre] - The multi-core processor of claim 1,

276. As explained above for Claim 1 in Ground 5, Naffziger and Allarey, alone or in combination, teach each and every limitation of Claim 1, rendering it obvious.

**2. 2[a] - wherein the interface block further comprises a first level shifter that is referenced to the second supply voltage and adapted to translate first logic levels associated with the first set of processor cores to second logic levels associated with the second set of processor cores for a first signal traveling from the first set of processor cores to the second set of processor cores.**

277. Naffziger teaches the use of a level shifter in the power control unit corresponding to the first core of Naffziger. Naffziger discloses that the power control unit 150 includes the voltage control unit 205 that "us[es] … level shifter circuity … configured to vary a supply voltage."

> Control logic 220 in the embodiment shown is coupled to both a voltage control unit 205 and a clock control unit 210. Control of the operating state of processor 100 may be accomplished by control logic 220 through manipulating one or both of a core supply voltage (V_core) or the frequency of a core clock signal (clk_core). Control logic 220 may alter the operating state of processor 100 between the low power state by directing voltage control unit 205 to change the core supply voltage provided to processor core 105. ***Voltage control unit 205 may be implemented using*** adjustable voltage regulator circuitry, ***level shifter circuitry***, or any other suitable circuitry ***configured to vary a supply voltage.*** An external voltage V may be received by voltage control unit 205 from a power supply via a motherboard to which processor 100 may be coupled. In one embodiment, control logic 220 may provide signals to voltage control unit 205 that cause processor 100 to enter the low power state by reducing the core supply voltage, and may cause processor 100 to enter the high performance state by increasing the core supply voltage.

Ex[1010] at [0032].

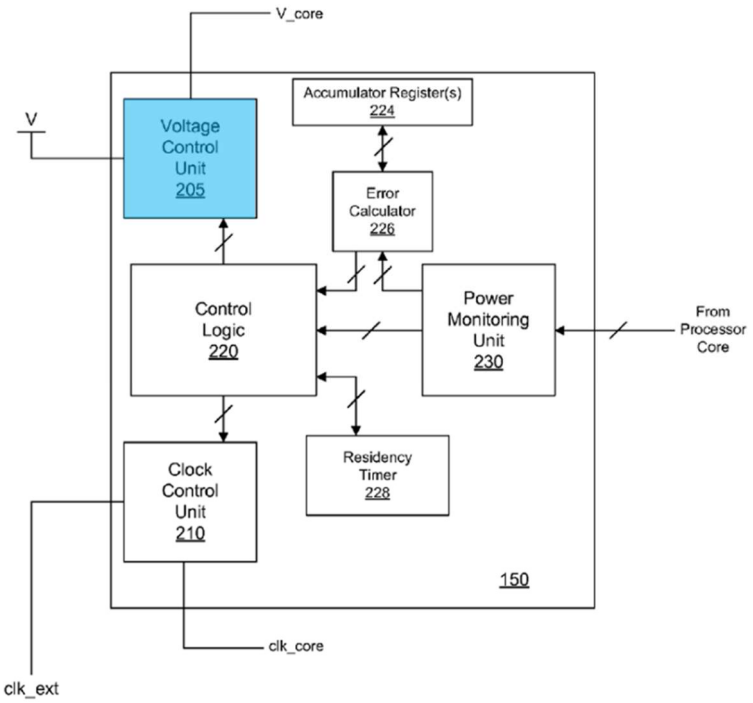


Fig. 2

*Fig. 2 of Naffziger, annotated*

278. As such, the power control unit corresponding to the first core 105A has a first level shifter and the power control unit corresponding to the second core 105B has a second level shifter.

279. Alternatively, the power control unit has "separate, independent core voltage supply planes" to separately control the first and second cores 105A and 105B. As discussed above for Claim 1[c1] in Ground 5, a PHOSITA reading Naffziger would have understood that each plane has a voltage control unit 205 comprising separate level shifter circuitry for independent control of voltage supply to each core. Accordingly, Naffziger discloses that the interface block comprises a first level shifter (highlighted in dashed blue block) and a second level shifter

(highlighted in solid blue block).

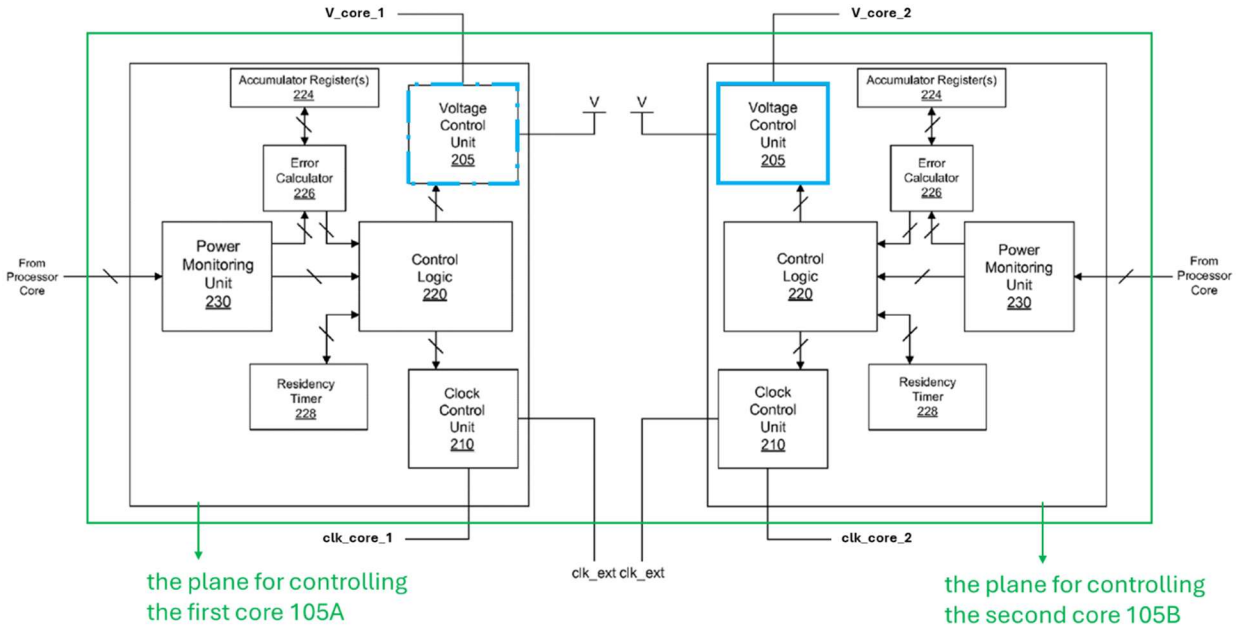


*Fig. 2 of Naffziger, modified to show two planes with a first level shifter (dashed blue block on the left) and a second level shifter (solid blue block on the right)*

280. Naffziger discloses that the first and second cores 105A and 105B operate in different states under different supply voltages.

> ***In balancing core power consumption with processor performance, power control unit 150 may be configured to cause processor 100 to operate in various power states. In one embodiment, power control unit 150 may cause processor 100 to alternate between operating in a high performance state and a low power state.*** In the high performance state, the core voltage and/or the clock frequency may be set by power control unit 150 to their maximum rated values. In the low power state, at least one (if not both) of the core voltage and core clock frequency may operate at a reduced value.

Ex[1010] at [0025].

> Turning now to FIG. 5, a block diagram of one embodiment of a processor having a power control unit and a plurality of processor cores is shown. In the embodiment shown, processor 500 is a dual core

> processor, having a first core 105A and a second core 105B. Processor 500 also includes a north bridge unit 110 and a power control unit 150, both of which are coupled to both processor cores 105A and 105B. ***Power control unit 150 is configured to monitor the power of both processor cores 105A and 105B, and configured to alternate operation of these cores between at least a high performance state and a low power state. As noted above, additional, intermediate operational states may be implemented. Power control unit 150 may accomplish the switching between operational states of processor cores 105A and 105B by changing the frequency of clock signals provided to the cores, changing supply voltages provided to the cores, or both.***
>
> In some embodiments, the operational state of cores 105A and 105B are in conjunction with each other. That is, both processor cores 105A and 105B are operated in the same state. However, ***embodiments are possible and contemplated wherein the operational states of processor cores 105A and 105B may be controlled independently of one another.*** For example, in such an embodiment, processor core **105**A may operate in the high performance state, while processor core **105**B operates in the low power state. ***Such a configuration may require that cores 105A and 105B have separate, independent core supply voltage planes and/or are configured to receive separate, independently controlled core clock signals.*** In addition, embodiments are possible and contemplated wherein power control unit 150 is configured to monitor the processing workloads of each processor core 105A and 105B, and may change operational states of the cores by changing their respective processing workloads.

*Id.* at [0053]-[0054].

281. When a signal is transmitted between cores operating with different voltage levels, the signal's logic levels may not match the expected input levels of the receiving domain. A PHOSITA would have understood that, to ensure proper interpretation and to prevent signal integrity issues or physical damages, it is necessary to translate the voltage level of the signal from the first set of cores to the

second set of cores. Level shifters were well-understood as a standard design choice to translate signals transmitted between circuits with different voltage levels. The routine application of a level shifter to handle inter-core signals is to translate the logic levels of one core to compatible logic levels of the other. That is, the first level shifter translates the logic levels of the first set of cores to compatible logic levels of the second set of cores for the signal transmitted from the first set of cores to the second set of cores; the second level shifter translates the logic levels of the second set of cores to compatible logic levels of the first set of cores for the signal transmitted from the first set of cores. A PHOSITA would have known that level shifters are a simple and routine design choice, with applicability in this context, based on the teaching in Naffziger and/or Allarey of the first and second level shifters in the interface block that facilitates communication between the first and second cores operating with different voltage levels.

### D. Dependent Claim 3

#### 1. 3[pre] - The multi-core processor or claim 1,

282. As explained above for Claim 1 in Ground 5, Naffziger and Allarey, alone or in combination, teach each and every limitation of Claim 1, rendering it obvious.

**2. 3[a] - wherein the interface block further comprises a second level shifter that is referenced to the first supply voltage and adapted to translate second logic levels associated with the second set of processor cores to first logic levels associated with the first set of processor cores for a second signal traveling from the second set of processor cores to the first set of processor cores.**

283. As discussed above for Claim 2[a] in Ground 5, Naffziger discloses that the interface block comprises a second level shifter that translates the logic levels of the second set of cores to compatible logic levels of the first set of cores for the signal transmitted from the first set of cores.

## E. Dependent Claim 5

### 1. 5[pre] - The multi-core processor of claim 1,

284. As explained above for Claim 1 in Ground 5, Naffziger and Allarey, alone or in combination, teach each and every limitation of Claim 1, rendering it obvious.

**2. 5[a] - wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor.**

285. As explained above for Claim 5[a] in Ground 1, Allarey teaches this claim limitation.

286. Naffziger discloses that "control logic 220" within power control unit 150 controls the operating states of the processor cores "through manipulating one or both of a core supply voltage (V_core) or the frequency of a core clock signal

(clk_core)." Ex[1010] at [0032]. Control logic 220 is coupled to both voltage control unit 205 and clock control unit 210. *Id.* Control logic 220 "provide[s] signals to voltage control unit 205" and "direct[s] voltage control unit 205 to change the core supply voltage provided to processor core 105." *Id.* Control logic 220 also provides signals to clock control unit 210 to "vary the frequency of the core clock signal." *Id.* at [0033].

> ***Control logic 220 in the embodiment shown is coupled to both a voltage control unit 205 and a clock control unit 210. Control of the operating state of processor 100 may be accomplished by control logic 220 through manipulating one or both of a core supply voltage (V_core) or the frequency of a core clock signal (clk_core***). ***Control logic 220 may alter the operating state of processor 100 between the low power state by directing voltage control unit 205 to change the core supply voltage provided to processor core 105.*** Voltage control unit 205 may be implemented using adjustable voltage regulator circuitry, level shifter circuitry, or any other suitable circuitry configured to vary a supply voltage. An external voltage V may be received by voltage control unit 205 from a power supply via a motherboard to which processor 100 may be coupled. ***In one embodiment, control logic 220 may provide signals to voltage control unit 205*** that cause processor 100 to enter the low power state by reducing the core supply voltage, and may cause processor 100 to enter the high performance state by increasing the core supply voltage.
>
> Clock control unit 210 may be implemented using a phase locked loop (PLL), a delay locked loop (DLL), or other type of circuitry that can be used for adjusting the frequency of a clock signal. In the embodiment shown, an external clock signal, clk_ext, is provided at a first frequency to clock control unit 210. A core clock signal, clk_core, is provided by clock control unit 210 to processor core 105. ***Clock control unit 210 is configured to vary the frequency of the core clock signal based on signals received from control logic 220.*** In one embodiment, control logic 220 may cause processor 100 to enter the low power state by directing clock control unit 210 to reduce the frequency of the core

> clock signal, and may cause the processor 100 to enter the high performance state by directing clock control unit to increase the frequency of the core clock signal.

*Id.* at [0032]-[0033].

287. As depicted below, control logic 220 (highlighted in red) is located in the periphery of Naffziger's multi-core processor.

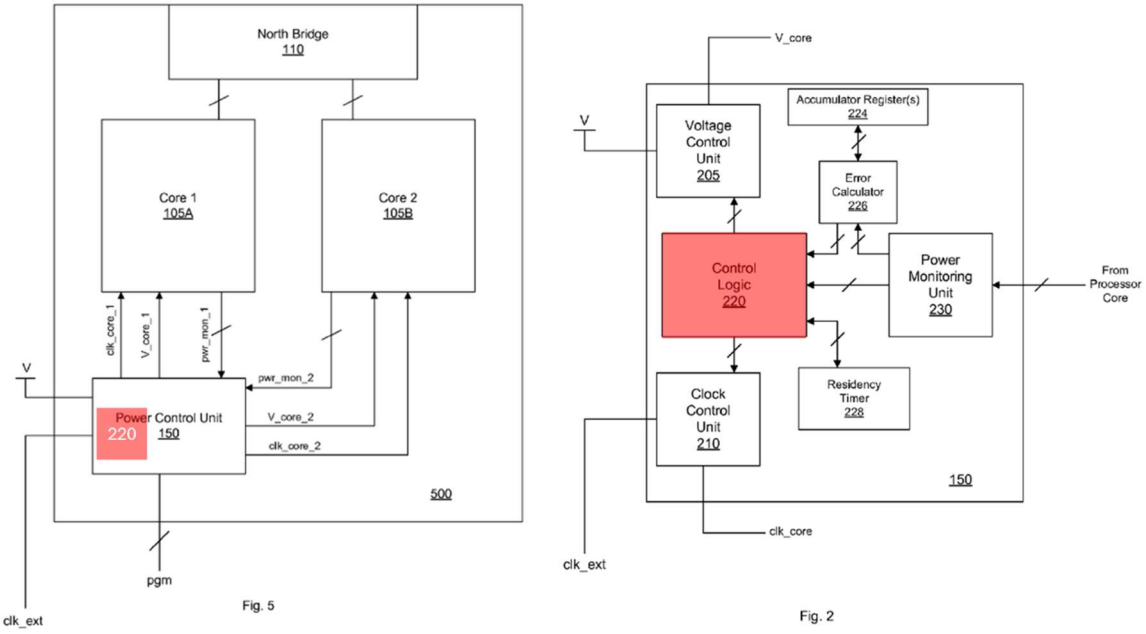


*Figs. 2 and 5 of Naffziger, annotated*

288. Furthermore, as explained above for Claim 5[a] in Ground 1, it was a known design option for a PHOSITA to place the control blocks in the periphery of the multi-core processor without affecting its function of providing control signals to the processor cores. A PHOSITA reading Naffziger would have found it obvious that control logic 220 could be placed in the periphery of the multi-core processor.

289. Accordingly, Naffziger and Allarey, alone or in combination, teach

Claim 5.

### F. Dependent Claim 8

#### 1. 8[pre] - The multi-core processor of claim 1,

290. As explained above for Claim 1 in Ground 5, Naffziger and Allarey, alone or in combination, teach each and every limitation of Claim 1, rendering it obvious.

#### 2. 8[a] - wherein the first set of processor cores are located in a first region of the multi-core processor, and the second set of processor cores are located in a second region of the multi-core processor.

291. As discussed above for Claim 8[a] in Ground 1, the '339 Patent discloses that a "region" is simply a spatial grouping of cores based on their physical location in the multi-core processor system. In Naffziger, the first set of processor cores (processor core 105A and its associated cores) are located in the physical space depicted below in red dotted box; the second set of processor cores (processor core 105B and its associated cores) are located in the physical space depicted below in green dotted box.

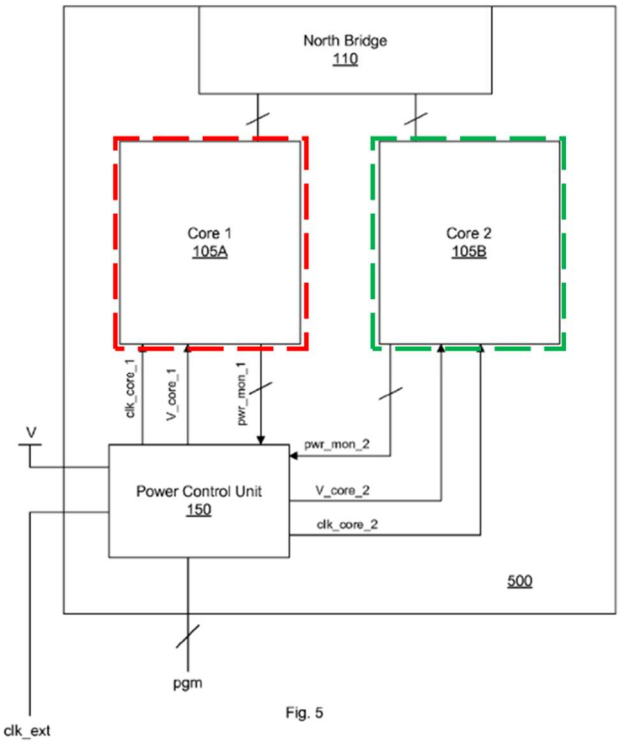


*Fig. 5 of Naffziger, annotated*

292. As such, Naffziger includes a first region corresponding to the physical location reflected by the red dotted box, and a second region corresponding to the physical location reflected by the green dotted box. The first and second regions depicted below are completely separate in physical layout and thus are non-overlapping.

293. For the same reasons described for Claim 8[a] in Ground 1, Allarey discloses this claim limitation.

294. Accordingly, Naffziger alone or Naffziger in view of Allarey renders Claim 8 obvious.

### G. Dependent Claim 9

#### 1. 9[pre] - The multi-core processor or claim 8,

295. As explained above for Claim 8 in Ground 5, Naffziger and Allarey, alone or in combination, teach each and every limitation of Claim 8, rendering it obvious.

#### 2. 9[a] - wherein the first region and the second region are overlapping regions of the multi-core processor.

296. For the same reasons described for Claim 9[a] in Ground 1, Allarey discloses this claim limitation.

297. Accordingly, Naffziger alone or Naffziger in view of Allarey renders Claim 9 obvious.

### H. Dependent Claim 10

#### 1. 10[pre] - The multi-core processor of claim 8,

298. As explained above for Claim 8 in Ground 5, Naffziger and Allarey, alone or in combination, teach each and every limitation of Claim 8, rendering it obvious.

#### 2. 10[a] - wherein the first region and the second region are non-overlapping regions of the multi-core processorS.

299. For the same reasons described above for Claim 8[a] in Ground 5, Naffziger discloses this claim limitation.

300. For the same reasons described for Claim 10[a] in Ground 1, Allarey discloses this claim limitation.

301. Accordingly, Naffziger alone or Naffziger in view of Allarey renders Claim 9 obvious.

## I. Dependent Claim 14

### 1. 14[pre] - The multi-core processor of claim 1,

302. As explained above for Claim 1, Naffziger and Allarey, alone or in combination, teach limitation of Claim 1, rendering it obvious.

### 2. 14[a] - wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a common region that is substantially central to the first set of processor cores and the second set of processor cores.

303. For the same reasons described for Claim 14[a] in Ground 1, Allarey discloses this claim limitation.

## J. Independent Claim 21

### 1. 21[pre] - A multi-core processor, comprising:

304. To the extent the preamble is limiting, for the same reasons described for Claim 1[Preamble] in Ground 5, Naffziger and Allarey, alone or in combination, disclose a multi-core processor.

### 2. 21[a1] - a first set of processor cores of the multi-core processor,

305. For the same reasons described for Claim 1[a1] in Ground 5, Naffziger and Allarey, alone or in combination, disclose a first set of processor cores of the multi-core processor.

**3. 21[a2] - wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage from a power control block and a first output clock signal from a first phase lock loop (PLL) having a first clock signal as input in a clock control block;**

306. For the same reasons described for Claim 1[a2] in Ground 5, Naffziger and Allarey, alone or in combination, disclose "wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage" and "a first output clock signal from a first phase lock loop (PLL) having a first clock signal as input."

307. As previously described and illustrated below in Figs. 2 and 5 of Naffziger, the first set of cores of Naffziger receives a first voltage supply "V_core_1" (shown by the blue line) from the voltage control unit 205 of the first power control unit associated with the first set of cores (highlighted in blue), and a first output clock signal "clk_core_1" (shown by the purple line) of a phase locked loop (PLL) in the clock control unit 210 of the first power control unit associated with the first set of cores (highlighted in purple). Ex[1010] at [0033].

> ***Clock control unit 210 may be implemented using a phase locked loop (PLL)***, a delay locked loop (DLL), or other type of circuitry that can be used ***for adjusting the frequency of a clock signal***. In the embodiment shown, an external clock signal, clk_ext, is provided at a first frequency to clock control unit 210. ***A core clock signal, clk_core, is provided by clock control unit 210 to processor core 105. Clock control unit 210 is configured to vary the frequency of the core clock signal based on signals received from control logic 220.*** In one embodiment, control logic 220 may cause processor 100 to enter the low power state by directing clock control unit 210 to reduce the frequency of the core

> clock signal, and may cause the processor 100 to enter the high performance state by directing clock control unit to increase the frequency of the core clock signal.

*Id.* at [0033].

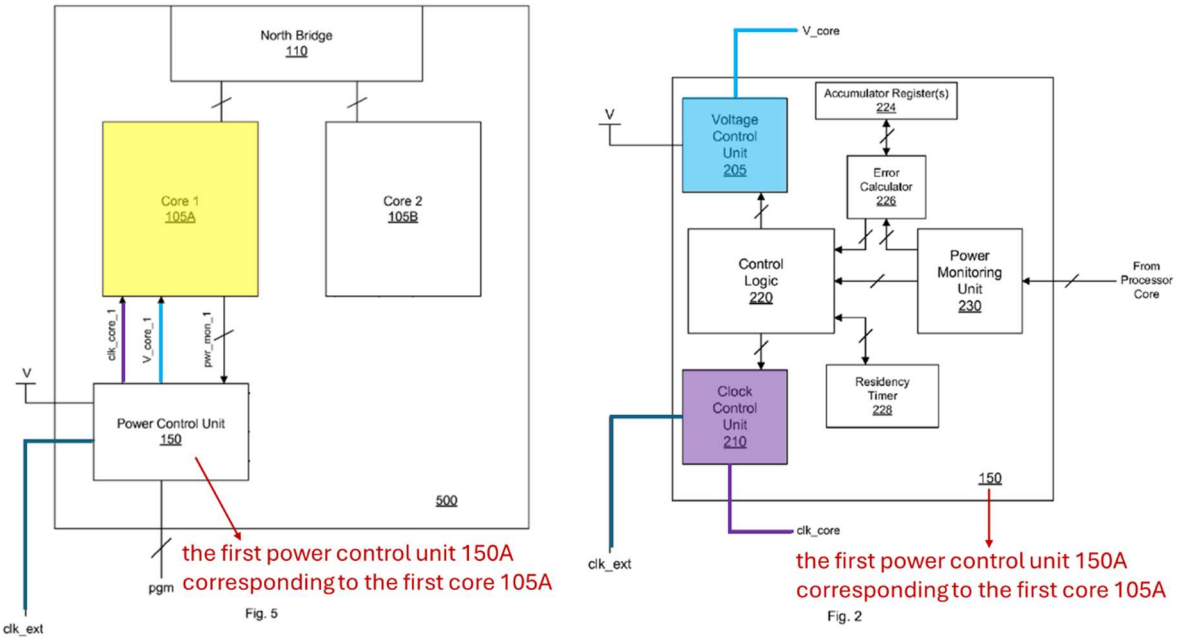


*Figs. 2 and 5 of Naffziger, annotated*

308. As previously described for Claim 1[b2] in Ground 5 and illustrated below in Figs. 2 and 5 of Naffziger, the second set of cores of Naffziger receives a second voltage supply "V_core_2" (shown by the blue line) from the voltage control unit 205 of the second power control unit associated with the second set of cores (highlighted in blue), and a second output clock signal "clk_core_2" (shown by the purple line) of a phase locked loop (PLL) in the clock control unit 210 of the second power control unit associated with the second set of core (highlighted in purple). Ex[1010] at [0033].

> ***Clock control unit 210 may be implemented using a phase locked loop (PLL)***, a delay locked loop (DLL), or other type of circuitry that can be used ***for adjusting the frequency of a clock signal***. In the embodiment shown, an external clock signal, clk_ext, is provided at a first frequency to clock control unit 210. ***A core clock signal, clk_core, is provided by clock control unit 210 to processor core 105. Clock control unit 210 is configured to vary the frequency of the core clock signal based on signals received from control logic 220.*** In one embodiment, control logic 220 may cause processor 100 to enter the low power state by directing clock control unit 210 to reduce the frequency of the core clock signal, and may cause the processor 100 to enter the high performance state by directing clock control unit to increase the frequency of the core clock signal.

*Id.* at [0033].

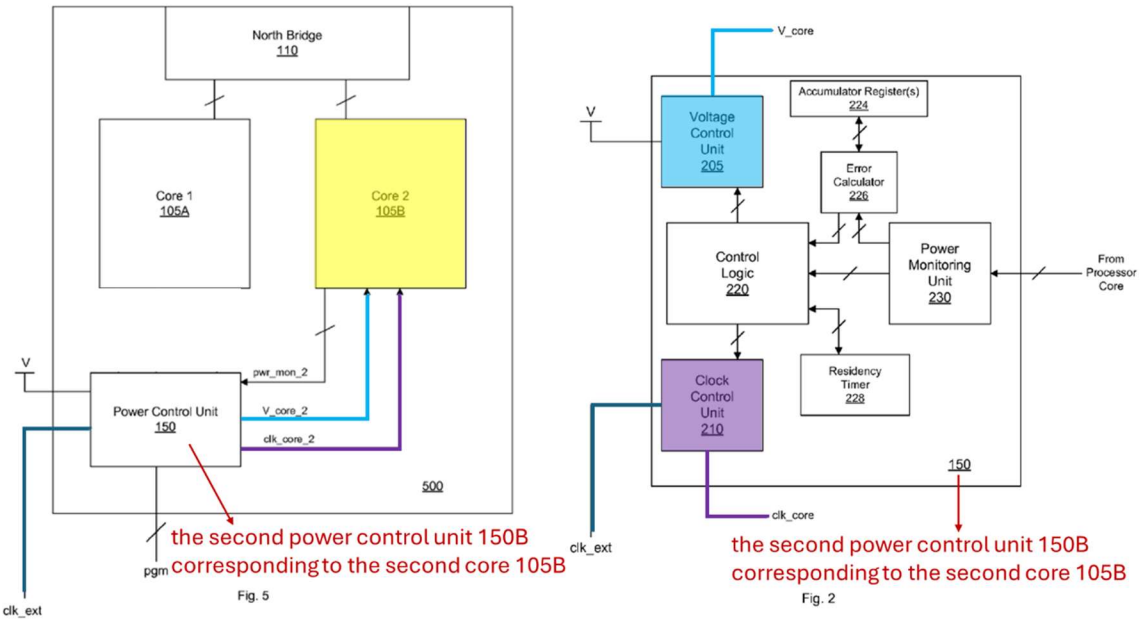


*Figs. 2 and 5 of Naffziger, annotated*

309. Naffziger thus discloses a power control block comprising the voltage control units 205 of the first and second power control units, and a clock control block comprising the clock control units 210 of the first and second power control

units.

310. Further, as previously described for Claim 1[c1] in Ground 5 and depicted below, Naffziger also teaches a single structure power control unit 150 that includes separate planes for controlling the first and second processor cores 105A and 105B independently. As depicted below, power control unit 150 includes the power control block comprising the voltage control units 205 of both planes (as shown in dashed blue block) and the clock control block comprising the clock control units 210 of both planes (as shown in dashed purple block).

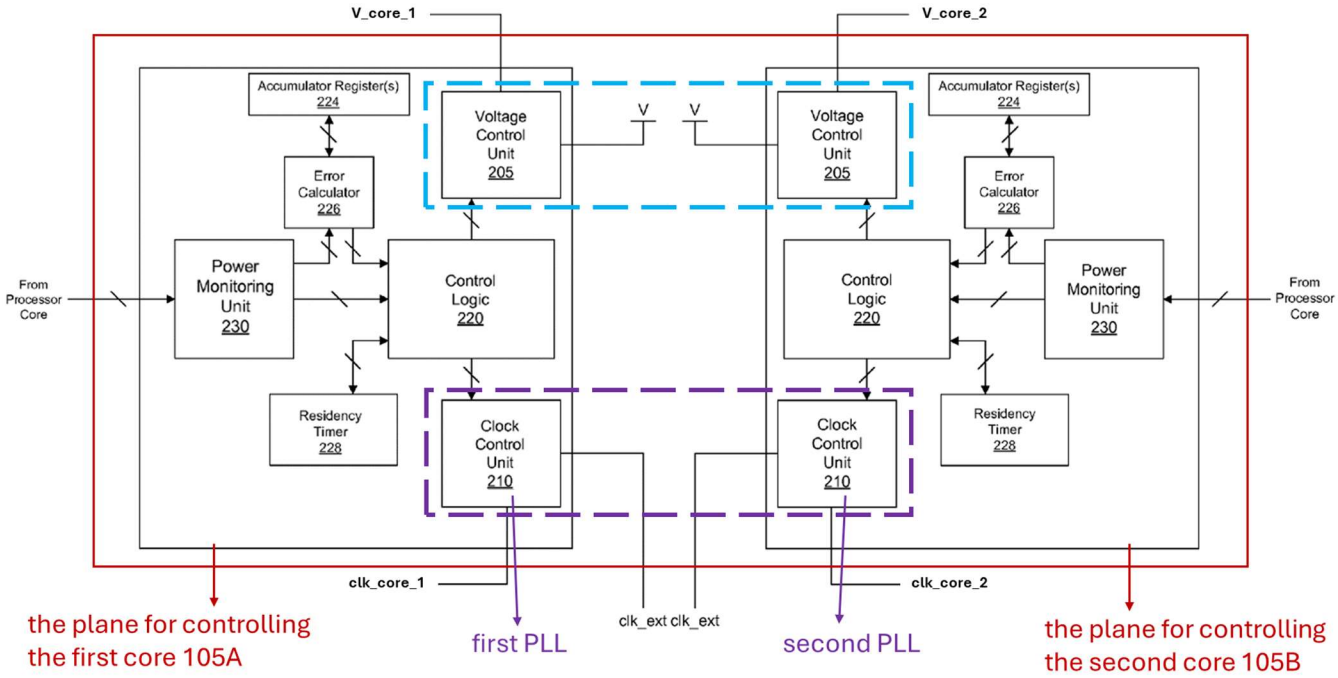


*Fig. 2 of Naffziger, modified to show separate planes*

311. Accordingly, Naffziger discloses this limitation.

312. For the same reasons described for Claim 21[a2] in Ground 1, Allarey also discloses this claim limitation.

**4. 21[b1] - a second set of processor cores of the multi-core processor,**

313. For the same reasons described for Claim 1[b1] in Ground 5 and as explained above for Claim 21[a1] in Ground 5, Naffziger discloses and/or renders obvious "a second set of processor cores of the multi-core processor." Furthermore, to the extent it is argued or found that Naffziger does not disclose this limitation, a PHOSITA would have been motivated to modify the architecture of Naffziger to include a first and second sets of cores in view of Allarey. Therefore, Naffziger alone or in combination with Allarey discloses "a second set of processor cores of the multi-core processor."

**5. 21[b2] - wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage from the power control block and a second output clock signal from a second PLL having a second clock signal as input in the clock control block,**

314. For the same reasons described for Claim 1[b2] in Ground 5, Naffziger discloses "wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage" and "a second output clock signal from a second phase lock loop (PLL) having a second clock signal as input."

315. As explained above for Claim 21[a2] in Ground 5, Naffziger further

discloses that the second set of cores receives a second voltage supply "V_core_2" from the power control block that includes the voltage control units 205 associated with the first and second cores 105A and 105B respectively. Naffziger also discloses that the second set of cores of Naffziger receives a second output clock signal "clk_core_2" from the second PLL (the PLL within the clock control unit 210 associated with the second core 105B). The second PLL is in the clock control block comprising the clock control units 210 associated with the first and second cores 105A and 105B respectively. Accordingly, Naffziger discloses this limitation.

316. For the same reasons described for Claim 21[a2] and [b2] in Ground 1, Allarey also discloses this claim limitation.

**6. 21[b3] - wherein the first supply voltage is independent from the second supply voltage, and**

317. For the same reasons described for Claim 1[b3] in Ground 5, Naffziger discloses this claim limitation.

318. For the same reasons described for Claim 1[b3] in Ground 1, Allarey also discloses this claim limitation.

**7. 21[b4] - the first clock signal is independent from the second clock signal; and**

319. For the same reasons described for Claim 1[b4] in Ground 5, Naffziger discloses this claim limitation.

320. For the same reasons described for Claim 1[b4] in Ground 1, Allarey

discloses this claim limitation.

**8. 21[c1] - an interface block coupled to the first set of processor cores and also coupled to the second set of processor cores,**

321. For the same reasons described for Claim 1[c1] in Ground 5, Naffziger discloses this claim limitation.

322. For the same reasons described for Claim 1[c1] in Ground 1, Allarey discloses this claim limitation.

**9. 1[c2] - wherein the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores.**

323. For the same reasons described for Claim 1[c2] in Ground 5, Naffziger alone or Naffziger in view of Allarey discloses "the interface block is configured to facilitate communication between the first set of processor cores and the second set of processor cores."

## XI. GROUND 6: CLAIM 4 IS OBVIOUS UNDER §103 OVER NAFFZIGER AND ALLAREY IN VIEW OF FLAUTNER

324. A PHOSITA would have understood that Naffziger (Ex[1010]) and Allarey (Ex[1006]), alone or in combination, further in view of Flautner (Ex[1007]) teach each and every limitation of Claim 4.

### A. Dependent Claim 4

**1. 4[pre] - The multi-core processor of claim 1,**

325. As explained above for Claim 1 in Ground 5, Naffziger and Allarey, alone or in combination, teach each and every limitation of Claim 1, rendering it

obvious.

**2. 4[a] - wherein the interface block further comprises a synchronizer configured to synchronize the first clock signal and the second clock signal for communication between one or more processor cores of the first set of processor cores and one or more processor cores of the second set of processor cores.**

326. As explained above for Claim 4 in Ground 2, Flautner discloses this claim element.

**B. Motivation to Combine Naffziger, Allarey and Flautner**

327. Naffziger and Allarey disclose complex multi-core processors that operate in different power states and require careful management of timing and synchronization to ensure efficient communication and operation. Flautner's synchronizer is specifically designed to handle the timing and voltage differences between processor cores, making it highly compatible with the needs of a multi-core system that operates under varying conditions.

> FIG. 4 illustrates a fourth embodiment. ***This embodiment is similar to that of FIG. 1 except that in this case the first processor core 40 and the second processor core 42 are asymmetrically controlled by the clock speed controller 44. More particularly, this asymmetric control allows clocks of different speeds to be simultaneously supplied to respective ones of the first processor core 40 and the second processor core 42. Thus, the first processor core 40 may be supplied with a relatively fast clock whilst the second processor core 42 is being supplied with a relatively slow clock.*** Accompanying these different speed clocks the integrated circuit upon which the first processor core 40 and the second processor core 42 are both formed may be split into multiple voltage domains with respective power controllers 46, 48. Thus, the second processor core 42 may have its clock speed reduced

> and its supply voltage lowered so as to reduce energy consumption by the second processor core 42 whilst the first processor core 40 maintains a high speed clock and a higher supply voltage needed for that high speed clock.
>
> When using different clock speeds and voltage levels in the first processor core 40 and the second processor core 42 it will be appreciated that ***a synchronisation module 50*** and a voltage level shifter 52 ***is provided between the first processor core 40 and the second processor core 42 to deal with clock synchronisation issues*** and the different supply voltage levels (voltage signalling levels) between the two domains.

Ex[1007] at [0047]-[0048].

328. Further, a PHOSITA would have understood that to facilitate communication between the first and second sets of cores of Naffziger and Allarey, the clock signals of the two cores would necessarily be required to be synchronized. When two cores operate at different clock speeds, the timing of signal transfers between the cores will not align. Synchronization of the timing of the clock signals is necessary to make sure the communication from one core is recognized and correctly interpreted by the other core. The synchronisation module 50 of Flautner thus synchronizes the clock signals of the first and second cores for communication between the first and second cores, as would be required in the systems of Naffziger and Allarey.

329. Naffziger and Allarey, further in view of Flautner, thus teach each and every limitation of Claim 4, rendering it obvious.

## XII. GROUND 7: CLAIM 6 IS OBVIOUS UNDER §103 OVER NAFFZIGER

**AND ALLAREY IN VIEW OF WOLFE, AND FURTHER IN VIEW OF KUMAR**

330. A PHOSITA would have understood that Naffziger (Ex[1010]) and Allarey (Ex[1006]), alone or in combination, in view of Wolfe (Ex[1008]) and/or Kumar (Ex[1009]) teaches or suggests each and every limitation of Claim 6 and therefore renders it obvious.

**A. Motivation to Combine Naffziger, Allarey and Wolfe**

331. A PHOSITA would have been motivated to combine Naffziger and Allarey. *See* Section X.A. A PHOSITA would have been further motivated to combine Naffziger and Allarey with Wolfe.

332. A PHOSITA would have recognized that Naffziger and Allarey focus on functional mechanisms for dynamic voltage and frequency regulation in a multi-core processor system. It would have been in the best interest of a PHOSITA to explore complementary configurations that optimize the physical layout of the multi-core processor system to support these mechanisms. For similar reasons as described above in Section VIII.A, Wolfe discloses a physical layout for a multi-core processor having independent control mechanisms for dynamic management of voltage and frequency levels that would have prompted a PHOSITA to consider combining with Naffziger and Allarey.

**B. Motivation to Combine Naffziger, Allarey and Kumar**

333. A PHOSITA would have been further motivated to combine Naffziger

and Allarey with Kumar. To improve the power management and inter-core voltage regulation of Naffziger's and Allarey's multi-core processor systems, a PHOSITA would have been inclined to look beyond the teachings of these references to identify improved configurations as part of the normal course of his/her own research. A PHOSITA would have been inclined to seek references which cover specific mechanisms for inter-core voltage regulation in a multi-core processor.

334. For similar reasons as described above in Section VIII.B, a PHOSITA would have been motivated to combine Naffziger and Allarey with Kumar because Kumar is directed to the same field of endeavor, *i.e.*, a power management system in the context of integrated circuits, and Kumar additionally discloses details for dynamically managing the respective voltage supplies within predetermined voltage relationships to optimize the overall power distribution and performance that can be easily implemented by a simple combination of known parts.

### C. Dependent Claim 6

#### 1. 6[pre] - The multi-core processor or claim 5,

335. As explained above for Claim 5 in Ground 5, Naffziger and Allarey, alone or in combination, teach each and every limitation of Claim 5, rendering it obvious.

#### 2. 6[a] - wherein the first set of processor cores is adjacent to the second set of processor cores, and

336. Naffziger and Allarey, alone or in combination, disclose "wherein the

first set of processor cores is adjacent to the second set of processor cores."

337. Naffziger discloses that processor cores 105A and 105B are placed adjacent to each other, as depicted below.

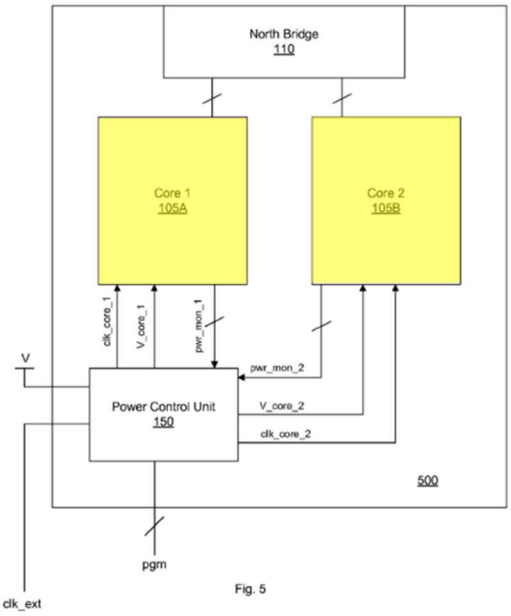


*Fig. 5 of Naffziger, annotated*

338. For the same reasons as discussed for Claim 6[a] in Ground 4, Allarey and Wolfe also disclose this claim limitation.

**3. 6[b] - the one or more control blocks are configured to select the first supply voltage and the second supply voltage to maintain a differential relationship between the first supply voltage and the second supply voltage.**

339. For the same reasons as discussed for Claim 6[b] in Ground 4, Kumar discloses this claim limitation.

340. Naffziger and Allarey, further in view of Kumar, thus teach each and every limitation of Claim 4, rendering it obvious.

## XIII. GROUND 8: CLAIM 11 IS OBVIOUS UNDER §103 OVER NAFFZIGER AND ALLAREY IN VIEW OF WOLFE

341. A PHOSITA would have understood that Naffziger in view of Allarey and further in view of Wolfe teaches or suggests each and every limitation of Claim 11 and therefore renders it obvious.

### A. 11[pre] - The multi-core processor or claim 8,

342. As explained above for Claim 8 in Ground 5, Naffziger and Allarey, alone or in combination, teach each and every limitation of Claim 8, rendering it obvious.

### B. 11[a] - wherein the first region corresponds to a first row of the multi-core processor, and wherein the second region corresponds to a second row of the multi-core processor.

343. For the same reasons as stated for Claim 11[a] in Ground 3, Wolfe discloses this claim limitation. Accordingly, Naffziger and/or Allarey in view of Wolfe renders Claim 11 obvious.

## XIV. SUMMARY OF CONCLUSIONS

344. Based on the foregoing analysis, the following chart summarizes my opinions regarding the invalidity of the ’339 patent:

| Ground | Opinion |
|---|---|
| 1 | Claims 1, 5, 8-10, 14 and 21 would have been obvious over Knoth in view of Allarey |
| 2 | Claims 2-4 would have been obvious over Knoth and Allarey in view of Flautner |

| Ground | Opinion |
|---|---|
| 3 | Claim 6 would have been obvious over Knoth and Allarey in view of Wolfe and further in view of Kumar |
| 4 | Claim 11 would have been obvious over Knoth and Allarey in view of Wolfe |
| 5 | Claims 1-3, 5, 8-10, 14 and 21 would have been obvious over Naffziger and Allarey |
| 6 | Claim 4 would have been obvious over Naffziger and Allarey in view of Flautner |
| 7 | Claim 6 would have been obvious over Naffziger and Allarey in view of Wolfe and further in view of Kumar |
| 8 | Claim 11 would have been obvious over Naffziger and Allarey in view of Wolfe |

## XV. SECONDARY CONSIDERATIONS

345. As discussed above, I have been informed that secondary considerations are a factor in the determination of obviousness. I have considered the current information related to secondary considerations of obviousness and determined that it has no impact on my conclusions. To the extent patent owner comes forward with any additional information, I will consider it.

## XVI. RIGHT TO SUPPLEMENT

346. This declaration represents the opinions on the '339 Patent that I have formed to date. I reserve the right to revise, supplement, and/or amend my opinions stated herein based on new information that becomes available to me and on my continuing analysis of the materials already provided. In connection with any arguments raised by patent owner, opinions by patent owner's expert(s), additional evidence and testimony, and/or judicial determinations, whether in this or a related proceeding, I reserve the right to supplement my opinions in the future to respond thereto.

I hereby declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct, that all statements made of my own knowledge are true, and that all statements made on information and belief are believed to be true. I understand that willful false statements are punishable by fine or imprisonment or both pursuant to 18 U.S.C. § 1004.

Dated: October 15, 2024

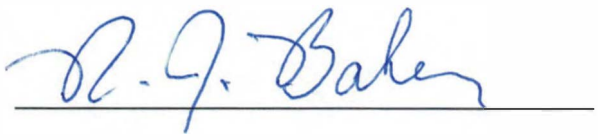

R. Jacob Baker, Ph.D., P.E.